

Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers *

Saeed Alaei [†]

Date: 2012-06-08 03:38:29 -0400 (Fri, 08 Jun 2012) [‡]

Abstract

For Bayesian combinatorial auctions, we present a general framework for approximately reducing the mechanism design problem for multiple buyers to single buyer sub-problems. Our framework can be applied to any setting which roughly satisfies the following assumptions: (i) buyers' types must be distributed independently (not necessarily identically), (ii) objective function must be linearly separable over the buyers, and (iii) except for the supply constraints, there should be no other inter-buyer constraints. Our framework is general in the sense that it makes no explicit assumption about buyers' valuations, type distributions, and single buyer constraints (e.g., budget, incentive compatibility, etc).

We present two generic multi buyer mechanisms which use single buyer mechanisms as black boxes; if an α -approximate single buyer mechanism can be constructed for each buyer, and if no buyer requires more than $\frac{1}{k}$ of all units of each item, then our generic multi buyer mechanisms are $\gamma_k \alpha$ -approximation of the optimal multi buyer mechanism, where γ_k is a constant which is at least $1 - \frac{1}{\sqrt{k+3}}$. Observe that γ_k is at least $\frac{1}{2}$ (for $k = 1$) and approaches 1 as $k \rightarrow \infty$. As a byproduct of our construction, we present a generalization of prophet inequalities. Furthermore, as applications of our framework, we present multi buyer mechanisms with improved approximation factor for several settings from the literature.

1 Introduction

The main challenge of stochastic optimization arises from the fact that all instances in the support of the distribution are relevant for the objective and this support is exponentially big in the size of problem. This paper aims to address this challenge by providing a general decomposition technique for assignment problems on independently distributed inputs where the objective is linearly separable over the inputs. The main challenge faced by such a decomposition approach is that the feasibility constraint of an assignment problem introduces correlation in the outcome of the optimal solution. In mechanism design problems, such constraints are typically the supply constraints. For example, when buyers are independent, a revenue maximizing seller with unlimited supply can decompose the problem over the buyers and optimize for each buyer independently. However, in the presence of supply constraints, a direct decomposition is not possible. Our decomposition

*A preliminary version of this paper was published in the proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2011).

[†]Department of Computer Science, University of Maryland, College Park, MD 20742, email: saeed@cs.umd.edu. Part of this work was done when the author was visiting Microsoft Research, New England. This work was partially supported by the NSF grant CCF-0728839

[‡]Id: bca.tex 201 2012-06-08 07:38:29Z saeed

technique can be roughly described as the following: (i) Construct a mechanism that satisfies the supply constraints only in expectation (ex-ante); the optimization problem for constructing such a mechanism can be fully decomposed over the set of buyers. (ii) Convert the mechanism from the previous step to another mechanism that satisfies the supply constraint at every instance.

We restrict our discussion to Bayesian combinatorial auctions. We are interested in mechanisms that allocate a set of heterogeneous items with limited supply to a set of buyers in order to maximize the expected value of a certain objective function which is linearly separable over the buyers (e.g., welfare, revenue, etc). The buyers' types are assumed to be distributed independently according to publicly known priors. We defer the formal statement of our assumptions to §2.

The following are the main challenges in designing mechanisms for multiple buyers.

- (I) The decisions made by the mechanism for different buyers should be coordinated because of supply constraints.
- (II) The decisions made by the mechanism for each buyer should be optimal (or approximately optimal).

Making coordinated optimal decisions for multiple buyers is challenging as it requires optimizing over the joint type space of all buyers, the size of which grows exponentially in the number of buyers. The second challenge is usually due to incentive compatibility (IC) constraints, specially in multi-dimensional settings where these constraints cannot be encoded compactly. In this paper, we mostly address the first challenge by providing a framework for approximately decomposing the mechanism design problem for multiple buyers to sub-problems dealing with each buyer individually.

Our framework can be summarized as follows. We start by relaxing the supply constraints, i.e., we consider the mechanisms for which only the ex-ante expected number of allocated units of each item is no more than the supply of that item. Note that “*ex-ante*” means that the expectation taken over all possible inputs (i.e., all possible types of the buyers). We show that the optimal mechanism for the relaxed problem can be constructed by independently running n single buyer mechanisms, where each single buyer mechanism is subject to an ex-ante probabilistic supply constraint. In particular, we show that if one can construct an α -approximate mechanism for each single buyer problem, then running these mechanisms simultaneously and independently yields an α -approximate mechanism for the relaxed multiple buyer problem. We then present two methods for converting the mechanism for the relaxed problem to a mechanism for the original problem while losing a small constant factor in the approximation. We present two generic multi buyer mechanisms that use the single buyer mechanisms from the previous step as blackboxes¹. In the first mechanism, we serve buyers sequentially by running, for each buyer, the corresponding single buyer mechanism from the previous step. However, we sometimes randomly preclude some of the items from the early buyers in order to ensure that late buyers get the same chance of being offered with those items; we ensure that the ex-ante expected probability of preclusion is equalized over all buyers, regardless of the order in which they are served (i.e., we simultaneously minimize the preclusion probability for all buyers). In the second mechanism, we run all of the single buyer mechanisms simultaneously and then modify the outcomes by deallocating some units of the over-allocated items at random while adjusting the payments respectively; we ensure that the ex-ante probability of deallocation is equalized among all units of each item and therefore simultaneously minimized for all buyers.

We also introduce a toy problem, *the magician's problem*, in §4, along with a near optimal solution for it, which is used as the main ingredient of our multi buyer mechanisms. As a byproduct,

¹Note that the single buyer mechanisms can be different for different buyers, e.g., to accommodate different classes of buyers.

we present improved generalized prophet inequalities for maximizing the sum of multiple choices.

As applications of our framework, in §6, we present mechanisms with improved approximation factor for several settings from the literature. For each setting we present a single buyer mechanism that satisfies the requirements of our framework, and can be plugged in one of our generic multi buyer mechanisms.

1.1 Related Work

In single dimensional settings, the related works from the CS literature are mostly focused on approximating the VCG mechanism for welfare maximization and/or approximating the Myerson's mechanism Myerson (1981) for revenue maximization (e.g., Bulow and Roberts (1989); Babaioff et al. (2006); Blumrosen and Holenstein (2008); Hartline and Roughgarden (2009); Dhangwatnotai et al. (2010); Chakraborty et al. (2010); Yan (2011)). Most of them consider mechanisms that have simple implementation and are computationally efficient. For welfare maximization in single dimensional settings, Hartline and Lucier (2010) gives a blackbox reduction from mechanism design to algorithmic design.

In multidimensional setting, for welfare maximization, Hartline et al. (2011) presents a blackbox reduction from mechanism design to algorithm design which subsumes the earlier work of Hartline and Lucier (2010). For revenue maximization, Chawla et al. (2010) presents several sequential posted pricing mechanisms for various settings with different types of matroid feasibility constraints. These mechanisms have simple implementation and approximate the revenue of the optimal mechanism. For unit-demand buyers whose valuations for the items are distributed according to product distributions, Chawla et al. (2010) present a sequential posted pricing mechanism that obtains in expectation at least $\frac{1}{6.75}$ -fraction of the revenue of the optimal posted pricing mechanism. In §6.2, we present an improved sequential posted pricing mechanism for this setting with an approximation factor of $\frac{1}{2}\gamma_k$ in which k is the number of units available of each item, and γ_k is a constant which is at least $1 - \frac{1}{\sqrt{k+3}}$. For combinatorial auctions with additive/correlated valuations with budget and demand constraints, Bhattacharya et al. (2010) presents all-pay $\frac{1}{4}$ -approximate BIC mechanisms for revenue maximization and a similar mechanism for welfare maximization. In subsection 6.4, we present an improved mechanism for this setting with an approximation factor of γ_k . Note that $\gamma_k k$ is at least $\frac{1}{2}$ and approaches 1 as $k \rightarrow \infty$. Bhattacharya et al. (2010) also presents sequential posted pricing mechanisms for the same setting, obtaining $O(1)$ approximation factors. For a similar setting, in §6.3, we present an improved sequential posted pricing mechanism with an approximation factor of $(1 - \frac{1}{e})\gamma_k$. Finally, Chawla et al. (2011) also considers various settings with hard budget constraints.

Prophet inequalities have been extensively studied in the past (e.g. Hill and Kertz (1992)). Prior to this work, the best known bound for the generalization to sum of k choices was $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ by Hajiaghayi et al. (2007). We improve this to $1 - \frac{1}{\sqrt{k+3}}$. Note that the current bound is tight for $k = 1$, and is useful even for small values of k .

2 Preliminaries

The framework of this paper is presented for combinatorial auctions, but it can be readily applied to Bayesian mechanism design in other contexts. We begin by defining the model and some notation.

Model. We consider the problem of selling m indivisible heterogeneous items to n buyers where there are k_j units of each item $j \in [m]$. All the relevant private information of each buyer $i \in [n]$

is represented by her type $t_i \in T_i$ where T_i is the type space of buyer i . Let $\mathbf{T} = T_1 \times \dots \times T_n$ be the space of all type profiles. The buyers' type profile $\mathbf{t} \in \mathbf{T}$ is distributed according to a publicly known prior \mathbf{D} . We use $X_{ij}(\mathbf{t})$ and $P_i(\mathbf{t})$ to denote the random variables² respectively for the allocation of item j to buyer i and the payment of buyer i , for type profile \mathbf{t} . For a mechanism M , the random variables for allocations and payments are denoted respectively by $X_{ij}^M(\mathbf{t})$ and $P_i^M(\mathbf{t})$. We are interested in computing a mechanism that (approximately) maximizes³ the expected value of a given objective function $\text{OBJ}(\mathbf{t}, X, P)$ where \mathbf{t} , X , and P respectively represent the types, the allocations, and the payments of all buyers. We are only interested in mechanisms which are within a given space \mathbb{M} of feasible mechanisms. Formally, we aim to compute a mechanism $M \in \mathbb{M}$ that (approximately) maximizes $\mathbf{E}_{\mathbf{t} \sim \mathbf{D}}[\text{OBJ}(\mathbf{t}, X^M(\mathbf{t}), P^M(\mathbf{t}))]$.

Assumptions. We make the following assumptions.

- (A1) **Independence.** The buyers' types must be distributed independently, i.e., $\mathbf{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ where \mathcal{D}_i is the distribution of types for buyer i . Note that for a buyer i who has multidimensional types, \mathcal{D}_i itself does not need to be a product distribution.
- (A2) **Linear Separability of Objective.** The objective function must be linearly separable over the buyers, i.e., $\text{OBJ}(t, X, P) = \sum_i \text{OBJ}_i(t_i, X_i, P_i)$ where t_i , X_i , and P_i respectively represent the type, the allocations, and the payment of buyer i .
- (A3) **Single-Unit Demands.** No buyer should ever need more than *one unit* of each item, i.e., $X_{ij}(\mathbf{t}) \in \{0, 1\}$ for all $\mathbf{t} \in \mathbf{T}$. This assumption is not necessary and is only to simplify the exposition; it can be removed as explained in §8.
- (A4) **Incentive Compatibility.** \mathbb{M} must be restricted to (Bayesian) incentive compatible mechanisms. By direct revelation principle this assumption is without loss of generality⁴,
- (A5) **Convexity.** \mathbb{M} must be a convex space. In other words, every convex combination of every two mechanisms from \mathbb{M} must itself be a mechanism in \mathbb{M} . A convex combination of two mechanisms $M, M' \in \mathbb{M}$ is another mechanism M'' which simply runs M with probability β and runs M' with probability $1 - \beta$, for some $\beta \in [0, 1]$. In particular, if \mathbb{M} is restricted to deterministic mechanisms, it is not convex; however if \mathbb{M} also includes mechanisms that randomize over deterministic mechanisms, then it is convex⁵.
- (A6) **Decomposability.** The set of constraints specifying \mathbb{M} must be decomposable to supply constraints (i.e., $\sum_i X_{ij} \leq k_j$, for each item j) and single buyer constraints (e.g., incentive compatibility, budget, etc). We define this assumption formally as follows. For any mechanism M , let $[[M]]_i$ be the single buyer mechanism perceived by buyer i , by simulating⁶ the other buyers according to their respective distributions \mathbf{D}_{-i} . Define $\mathbb{M}_i = \{[[M]]_i | M \in \mathbb{M}\}$ to be the space of all feasible single buyer mechanisms for buyer i . The decomposability assumption

²Note that these random variables are often correlated. Furthermore, for a deterministic mechanism, these variables take deterministic values as a function of \mathbf{t} .

³All of the results of this paper can be applied to minimization problems by simply maximizing the negation of the objective function.

⁴It is WLOG, given that we are only interested in mechanisms that have Bayes-Nash equilibria.

⁵For an example of a randomized non-convex space of mechanisms, consider the space of mechanisms where the expected payment of every type must be either less than \$2 or more than \$4.

⁶The single buyer mechanism induced on buyer i can be obtained by simulating all buyers other than i by drawing a random \mathbf{t}_{-i} from \mathbf{D}_{-i} and running M on buyer i and the $n - 1$ simulated buyers with types \mathbf{t}_{-i} ; note that this is a single-buyer mechanism because the simulated buyers are just part of the mechanism.

requires that for any arbitrary mechanism M the following holds: if M satisfies the supply constraints and also $[[M]]_i \in \mathbb{M}_i$ (for all buyers i), then it must be that $M \in \mathbb{M}$.

We shall clarify the last assumption by giving an example. Suppose \mathbb{M} is the space of all buyer specific item pricing mechanisms, then \mathbb{M} satisfies the last assumption. On the other hand, if \mathbb{M} is the space of mechanisms that offer the same set of prices to every buyer, it does not satisfy the decomposability assumption, because there is an implicit inter-buyer constraint that the same prices should be offered to different buyers.

Throughout the rest of this paper, we often omit the range of the sums whenever the range is clear from the context (e.g., \sum_i means $\sum_{i \in [n]}$, and \sum_j means $\sum_{j \in [m]}$).

Multi buyer problem. Formally, the multi buyer problem is to find a mechanism M which is a solution to the following program.

$$\begin{aligned} & \text{maximize} && \sum_i \mathbf{E}_{\mathbf{t} \sim \mathbf{D}} [\text{OBJ}_i(t_i, X_i^M(\mathbf{t}), P_i^M(\mathbf{t}))] && (OPT) \\ & \text{subject to} && \sum_i X_{ij}^M(\mathbf{t}) \leq k_j, && \forall \mathbf{t} \in \mathbf{T}, \forall j \in [m] \\ & && [[M]]_i \in \mathbb{M}_i, && \forall i \in [n] \end{aligned}$$

Observe that, in the absence of the first set of constraints, we could optimize the mechanism for each buyer independently. This observation is the key to our multi to single buyer decomposition, which allows us to approximately decompose/reduce the multi buyer problem to single buyer problems. A mechanism M is an α -approximation of the optimal mechanism if it is a feasible mechanism for the above program and obtains at least α -fraction of the optimal objective value of the program.

Ex ante allocation rule. For a multi-dimensional mechanism M , the *ex ante allocation rule* is a vector $x \in [0, 1]^{n \times m}$ in which $x_{ij} = \mathbf{E}_{\mathbf{t} \sim \mathbf{D}}[X_{ij}^M]$ is the expected probability of allocating a unit of item j to buyer i , where the expectation is taken over all possible type profiles. Note that for any feasible mechanism M , by linearity of expectation, the ex ante allocation rule satisfies $\sum_i x_{ij} \leq k_j$, for every item j .

Single buyer problem. The *single buyer problem*, for buyer i , is to compute an optimal single buyer mechanism and its expected objective value, subject to a given upper bound $\bar{x}_i \in [0, 1]^m$ on the ex ante allocation rule; in other words, the single buyer mechanism may not allocate a unit of item j to buyer i with an expected probability of more than \bar{x}_{ij} , where the expectation is taken over $t_i \sim \mathcal{D}_i$. Formally, the single buyer problem is to compute the optimal value of the following program along with a corresponding solution (i.e., the optimal M_i), for a given \bar{x}_i .

$$\begin{aligned} & \text{maximize} && \mathbf{E}_{t_i \sim \mathcal{D}_i} [\text{OBJ}_i(t_i, X_i^{M_i}(t_i), P_i^{M_i}(t_i))] && (OPT_i) \\ & \text{subject to} && \mathbf{E}_{t_i \sim \mathcal{D}_i} [X_{ij}^{M_i}(t_i)] \leq \bar{x}_{ij}, && \forall j \in [m] \\ & && M_i \in \mathbb{M}_i, \end{aligned}$$

We typically denote an optimal single buyer mechanism for buyer i , subject to a given \bar{x}_i , by $M_i(\bar{x}_i)$, and denote its expected objective value (i.e., the optimal value of the above program as a function of \bar{x}_i) by $R_i(\bar{x}_i)$. Later, we prove that $R_i(\bar{x}_i)$, which we refer to as the *optimal benchmark* for buyer

i , is a concave function of \bar{x}_i . In the case of approximation, we say that a single buyer mechanism M_i together with a concave benchmark R_i provide an α -approximation of the optimal single buyer mechanism/optimal benchmark, if the expected objective value of $M_i(\bar{x}_i)$ is at least $\alpha R_i(\bar{x}_i)$ and if $R_i(\bar{x}_i)$ is an upper bound on the optimal benchmark, for every \bar{x}_i .

To make the exposition more concrete, consider the following single buyer problem as an example. Suppose there is only one type of item (i.e., $m = 1$) and the objective is to maximize the expected revenue⁷. Suppose buyer i 's valuation is drawn from a regular distribution with CDF, F_i . The optimal single buyer mechanism for i , subject to $\bar{x}_i \in [0, 1]$, is a deterministic mechanism which offers the item at some fixed price, while ensuring that the probability of sale (i.e., the probability of buyer i 's valuation being above the offered price) is no more than \bar{x}_i . In particular, the optimal benchmark $R_i(\bar{x}_i)$ is the optimal value of the following convex program as a function of \bar{x}_i .

$$\begin{aligned} & \text{maximize} && x_i F_i^{-1}(1 - x_i) \\ & \text{subject to} && x_i \leq \bar{x}_i \\ & && x_i \in [0, 1] \end{aligned}$$

Furthermore, the optimal single buyer mechanism offers the item at the price $F_i^{-1}(1 - x_i)$ where x_i is the optimal assignment for the above convex program. Note that, for a regular distribution, $x_i F_i^{-1}(1 - x_i)$ is concave in x_i , so the above program is a convex program.

3 Decomposition via Ex ante Allocation Rule

In this section we present general methods for approximately decomposing/reducing the multi buyer problem to single buyer problems. Recall that a single buyer problem is to compute the optimal single buyer mechanism $M_i(\bar{x}_i)$ and its expected objective value $R_i(\bar{x}_i)$ (i.e., the optimal benchmark), subject to an upper bound \bar{x}_i on the ex ante allocation rule. We present two methods for constructing an approximately optimal multi buyer mechanism, using M_i and R_i as black box. Furthermore, we show that if we can only compute an α -approximation of the optimal single buyer mechanism/optimal benchmark for each buyer i , then the factor α simply carries over to the approximation factor of the final multi buyer mechanism.

Multi buyer benchmark. We start by showing that the optimal value of the following convex program gives an upper bound on the expected objective value of the optimal multi buyer mechanism.

$$\begin{aligned} & \text{maximize} && \sum_i R_i(\bar{x}_i) && (\overline{OPT}) \\ & \text{subject to} && \sum_i \bar{x}_{ij} \leq k_j, && \forall j \in [m] \\ & && \bar{x}_{ij} \in [0, 1], \end{aligned}$$

We first show that the above program is indeed a convex program.

Theorem 1. *The optimal benchmarks R_i are always concave.*

⁷The optimal multi buyer mechanism for this setting is given by Myerson (1981); yet we consider this setting to keep the example simple and intuitive.

Proof. We prove this for an arbitrary buyer i . Let M_i and R_i denote the optimal single buyer mechanism and the optimal benchmark for buyer i . To show that R_i is concave, it is enough to show that for any $\bar{x}_i, \bar{x}'_i \in [0, 1]^m$ and any $\beta \in [0, 1]$, the following inequality holds.

$$R_i(\beta\bar{x}_i + (1 - \beta)\bar{x}'_i) \geq \beta R_i(\bar{x}_i) + (1 - \beta)R_i(\bar{x}'_i)$$

Consider the single buyer mechanism M'' that works as follows: M'' runs $M_i(\bar{x}_i)$ with probability β and runs $M_i(\bar{x}'_i)$ with probability $1 - \beta$. Note that \mathbb{M}_i is a convex space (this follows from A5 and A6), therefore $M'' \in \mathbb{M}_i$. Observe that by linearity of expectation, the ex ante allocation rule of M'' is no more than $\beta\bar{x}_i + (1 - \beta)\bar{x}'_i$ and the expected objective value of M'' is exactly $\beta R_i(\bar{x}_i) + (1 - \beta)R_i(\bar{x}'_i)$. So the expected objective value of the optimal single buyer mechanism, subject to $\beta\bar{x}_i + (1 - \beta)\bar{x}'_i$, may only be higher. That implies $R_i(\beta\bar{x}_i + (1 - \beta)\bar{x}'_i) \geq \beta R_i(\bar{x}_i) + (1 - \beta)R_i(\bar{x}'_i)$ which proves the claim. \square

Theorem 2. *The optimal value of the convex program (\overline{OPT}) is an upper bound on the expected objective value of the optimal multi buyer mechanism.*

Proof. Let M^* be an optimal multi buyer mechanism. Let x^* denote the ex ante allocation rule corresponding to M^* , i.e., $x^*_{ij} = \mathbf{E}_{\mathbf{t} \sim \mathbf{D}}[X_{ij}^{M^*}]$. Observe that x^* is a feasible assignment for the convex program and yields an objective value of $\sum_i R_i(x^*_i)$ which is upper bounded by the optimal value of the convex program. So to prove the theorem it is enough to show that the contribution of each buyer i to the expected objective value of M^* is upper bounded by $R_i(x^*_i)$. Consider $M_i^* = [[M^*]]_i$, i.e, the single buyer mechanism induced by M^* on buyer i . M_i^* can be obtained by simply running M^* on buyer i and simulating the other $n - 1$ buyers with random types $\mathbf{t}_{-i} \sim \mathbf{D}_{-i}$; Observe that M_i^* is a feasible single buyer mechanism subject to x^*_i and obtains the same expected objective value as M^* from buyer i , so the expected objective value of the optimal single buyer mechanism subject to x^*_i could only be higher. \square

Constructing multi buyer mechanisms. Theorem 2 suggests that by computing an optimal assignment of \bar{x} for the convex program (\overline{OPT}) and running the single buyer mechanism $M_i(\bar{x}_i)$ for each buyer i , one might obtain a reasonable multi buyer mechanism; however such a multi buyer mechanism would only satisfy the supply constraints in expectation; in other words, there is a good chance that some items are over allocated with a non-zero probability. We present two generic multi buyer mechanisms for combining the single buyer mechanisms and resolving the conflicts in the allocations in such a way that would ensure the supply constraints are met at every instance and not just in expectation. In both approaches we first solve the convex program (\overline{OPT}) to compute the optimal \bar{x} . The high level idea of each mechanism is explained below.

1. **Pre-Rounding.** This mechanism serves the buyers sequentially (arbitrary order); for each buyer i , it selects a subset S_i of available items and runs the single buyer mechanism $M_i(\bar{x}_i[S_i])$, where $\bar{x}_i[S_i]$ denotes the vector resulting from \bar{x}_i by zeroing the entries corresponding to items not in S_i . In particular, this mechanism sometimes precludes some of the available items from early buyers to make them available to late buyers. We show that if there are at least k units of each item, then S_i includes item j with probability at least $1 - \frac{1}{\sqrt{k+3}}$, for each buyer i and each item j .
2. **Post-Rounding.** This mechanism runs $M_i(\bar{x}_i)$ for all buyers i simultaneously and independently. It then modifies the outcomes by deallocating the over allocated items at random in such a way that the probability of deallocation observed by all buyers are equal, and therefore minimized over all buyers. The payments are adjusted respectively. We show that if there

are at least k units of each item, every allocation is preserved with probability $1 - \frac{1}{\sqrt{k+3}}$ from the perspective of the corresponding buyer.

We will explain the above mechanisms in more detail in §5 and present some technical assumptions that are sufficient to ensure that they retain at least $1 - \frac{1}{\sqrt{k+3}}$ fraction of the expected objective value of each $M_i \langle \bar{x}_i \rangle$.

Approximately optimal single buyer mechanisms. Throughout the above discussion, we assumed that we can compute the optimal single buyer mechanisms and the corresponding optimal benchmarks. However, it is likely that we can only compute an approximation of them. Suppose for each buyer i , M_i and R_i , instead of being optimal, only provide an α -approximation of the optimal single buyer mechanism/optimal benchmark, and suppose R_i is concave; then we can still use M_i and R_i in the above construction, but the final approximation factor will be multiplied by α .

Main result. The following informal theorem summarizes the main result of this paper. The formal statement of this result can be found in Theorem 7 and Theorem 8.

Theorem 3 (Market Expansion). *If for each buyer $i \in [n]$, an α -approximate single buyer mechanism M_i and a corresponding concave benchmark R_i can be constructed in polynomial time, then, with some further assumptions (explained later), a multi buyer mechanism $M \in \mathbb{M}$ can be constructed in polynomial time by using M_i as building blocks, such that M is $\gamma_k \alpha$ -approximation of the the optimal multi buyer mechanism in \mathbb{M} , where $k = \min_j k_j$ and γ_k is a constant which is at least $1 - \frac{1}{\sqrt{k+3}}$.*

In order to explain the construction of the multi buyer mechanism, we shall first describe the magician's problem and its solution, which is used in both pre-rounding and post-rounding for equalizing the expected probabilities of preclusion/deallocation over all buyers.

4 The Magician's Problem

In this section, we present an abstract online stochastic toy problem and a near-optimal solution for it which provides the main ingredient for combining single buyer mechanisms to form multi buyer mechanisms; it is also used to prove a generalized prophet inequality.

Definition 1 (The Magician's Problem). *A magician is presented with a series of boxes one by one, in an online fashion. There is a prize hidden in one of the boxes. The magician has k magic wands that can be used to open the boxes. If a wand is used on box i , it opens, but with a probability of at most x_i , which written on the box, the wand breaks. The magician wishes to maximize the probability of obtaining the prize, but unfortunately the sequence of boxes, the written probabilities, and the box in which the prize is hidden are arranged by a villain, and the magician has no prior information about them (not even the number of the boxes). However, it is guaranteed that $\sum_i x_i \leq k$, and that the villain has to prepare the sequence of boxes in advance (i.e., cannot make any changes once the process has started).*

The magician could fail to open a box either because (a) he might choose to skip the box, or (b) he might run out of wands before getting to the box. Note that once the magician fixes his strategy, the best strategy for the villain is to put the prize in the box that has the lowest ex ante probability of being opened, based on the magician's strategy. Therefore, in order for the magician

to obtain the prize with a probability of at least γ , he has to devise a strategy that guarantees an ex ante probability of at least γ for opening each box. Notice that allowing the prize to be split among multiple boxes does not affect the problem. It is easy to show the following strategy ensures an ex ante probability of at least $\frac{1}{4}$ for opening each box: for each box randomize and use a wand with probability $\frac{1}{2}$. But can we do better? We present an algorithm parameterized by a probability $\gamma \in [0, 1]$ which guarantees a minimum ex-ante probability of γ for opening each box while trying to minimize the number of wands broken. In Theorem 4, we show that for $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$ this algorithm never requires more than k wands.

Definition 2 (γ -Conservative Magician). *The magician adaptively computes a sequence of thresholds $\theta_1, \theta_2, \dots \in \mathbb{N}_0$ and makes a decision about each box as follows: let W_i denote the number of wands broken prior to seeing the i^{th} box; the magician makes a decision about box i by comparing W_i against θ_i ; if $W_i < \theta_i$, it opens the box; if $W_i > \theta_i$, it does not open the box; and if $W_i = \theta_i$, it randomizes and opens the box with some probability (to be defined). The magician chooses the smallest threshold θ_i for which $\Pr[W_i \leq \theta_i] \geq \gamma$ where the probability is computed ex ante (i.e., not conditioned on past broken wands). Note that γ is a parameter that is given. Let $F_{W_i}(\ell) = \Pr[W_i \leq \ell]$ denote the ex ante CDF of random variable W_i , and let S_i be the indicator random variable which is 1 iff the magician opens the box i . Formally, the probability with which the magician should open box i condition on W_i is computed as follows.*

$$\Pr[S_i = 1 | W_i] = \begin{cases} 1 & W_i < \theta_i \\ (\gamma - F_{W_i}(\theta_i - 1)) / (F_{W_i}(\theta_i) - F_{W_i}(\theta_i - 1)) & W_i = \theta_i \\ 0 & W_i > \theta_i \end{cases} \quad (s)$$

$$\theta_i = \min\{\ell | F_{W_i}(\ell) \geq \gamma\} \quad (\theta)$$

Observe that θ_i is in fact computed before seeing box i itself.

Define $s_i^\ell = \Pr[S_i = 1 | W_i = \ell]$; the CDF of W_{i+1} can be computed from the CDF of W_i and x_i as follows (assume x_i is the exact probability of breaking a wand for box i).

$$F_{W_{i+1}}(\ell) = \begin{cases} s_i^\ell x_i F_{W_i}(\ell - 1) + (1 - s_i^\ell x_i) F_{W_i}(\ell) & i \geq 1, \ell \geq 0 \\ 1 & i = 0, \ell \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (F_W)$$

Furthermore, if each x_i is just an upper bound on the probability of breaking a wand on box i , then the above definition of $F_{W_i}(\cdot)$ stochastically dominates the actual CDF of W_i , and the magician opens each box with a probability of at least γ .

In order to prove that a γ -conservative magician does not fail for a given choice of γ , we must show that the thresholds θ_i are no more than $k - 1$. The following theorem states a condition on γ that is sufficient to guarantee that $\theta_i \leq k - 1$ for all i .

Theorem 4 (γ -Conservative Magician). *For any $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$, a γ -conservative magician with k wands opens each box with an ex ante probability of at least γ . Furthermore, if x_i is the exact probability (not just an upper bound) of breaking a wand on box i for each i , then each box is opened with an ex ante probability exactly γ ⁸*

Proof. See §7. □

⁸In particular the fact that the probability of the event of breaking a wand for the i^{th} box is exactly x_i , conditioned on any sequence of prior events, implies that these events are independent for different boxes.

Definition 3 (γ_k). We define γ_k to be the largest probability such that for any $k' \geq k$ and any instance of the magician's problem with k' wands, the thresholds computed by a γ_k -conservative magician are less than k' . In other words, γ_k is the optimal choice of γ which works for all instances with $k' \geq k$ wands. By Theorem 4, we know that γ_k must be⁹ at least $1 - \frac{1}{\sqrt{k+3}}$.

Observe that γ_k is a non-decreasing function in k which is at least $\frac{1}{2}$ (when $k = 1$) and approaches 1 as $k \rightarrow \infty$. The next theorem shows that the lower bound of $1 - \frac{1}{\sqrt{k+3}}$ on γ_k cannot be far from the optimal.

Theorem 5 (Hardness of Magician's Problem). For any $\epsilon > 0$, it is not possible to guarantee an *ex ante* probability of $1 - \frac{k^k}{e^k k!} + \epsilon$ for opening each box (i.e., no magician can guarantee it). Note that $1 - \frac{k^k}{e^k k!} \approx 1 - \frac{1}{\sqrt{2\pi k}}$ by Stirling's approximation.

Proof. See §A. □

Prophet Inequalities. We prove a generalization of prophet inequalities by a direct reduction to the magician's problem.

Definition 4 (k -Choice Sum). A sequence of n non-negative random numbers V_1, \dots, V_n are drawn from arbitrary distributions F_1, \dots, F_n one by one in an arbitrary order. A gambler observes the process and may select k of the random numbers, with the goal of maximizing the sum of the selected ones; a random number may only be selected at the time it is drawn, and it cannot be unselected later. The gambler knows all the distributions in advance, and observes from which distribution the current number is drawn, but not the order in which the future numbers are drawn. On the other hand, a prophet knows all the actual draws in advance, so he chooses the k highest draws. We assume that the order in which the random numbers are drawn is fixed in advance (i.e., may not change based on the decisions of the gambler).

Hajiaghayi et al. (2007) proved that there is a strategy for the gambler that guarantees in expectation at least $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ fraction of the payoff of the prophet, using a non-decreasing sequence of k stopping rules (thresholds)¹⁰. Next, we construct a gambler that obtains in expectation at least γ_k fraction of the prophet's payoff, using a γ_k -conservative magician as a black box. Note that $\gamma_k \geq 1 - \frac{1}{\sqrt{k+3}}$. This gambler uses only a single threshold. However, he may skip some of the random variables at random.

Theorem 6 (Prophet Inequalities – k -Choice Sum). The following strategy ensures that the gambler obtains at least γ_k fraction of payoff of the prophet in expectation.¹¹

- Find a threshold τ such that $\sum_i \Pr[V_i > \tau] = k$ (e.g., by doing a binary search on τ).
- Use a γ_k -conservative magician with k wands. Upon seeing each V_i , create a box and write $x_i = \Pr[V_i > \tau]$ on it and present it to the magician. If the magician chooses to open the box and also $V_i > \tau$, then select V_i and break the magician's wand, otherwise skip V_i .

⁹Because for any $k' \geq k$ obviously $1 - \frac{1}{\sqrt{k+3}} \leq 1 - \frac{1}{\sqrt{k'+3}}$.

¹⁰A gambler with stopping rules τ_1, \dots, τ_k works as follows. Upon seeing V_i , he selects it iff $V_i \geq \tau_{j+1}$ where j is the number of random draws selected so far.

¹¹To simplify the exposition we assume that the distributions do not have point masses. The result holds with slight modifications if we allow point masses.

Proof. First, we compute an upper bound on the expected payoff of the prophet. Let x_i be the ex ante probability (i.e., before any random number is drawn) that the prophet chooses V_i (i.e. the probability that V_i is among the k highest draws). Let $u_i(x_i)$ denote the maximum possible contribution of the random variable V_i to the expected payoff of the prophet if V_i is selected with an ex ante probability x_i . Note that $u_i(x_i)$ is equal to the expected value of V_i conditioned on being above the $1 - x_i$ quantile, multiplied by the probability of V_i being above that quantile. Assuming $F_i(\cdot)$ and $f_i(\cdot)$ denote the CDF and PDF of V_i , we can write $u_i(x_i) = \int_{F_i^{-1}(1-x_i)}^{\infty} v f_i(v) dv$. By changing the integration variable and applying the chain rule we get $u_i(x_i) = \int_0^{x_i} F_i^{-1}(1-x) dx$. Observe that $\frac{d}{dx_i} u_i(x_i) = F_i^{-1}(1-x_i)$ is a non-increasing function, so $u_i(x_i)$ is a concave function. Furthermore, $\sum_i x_i \leq k$ because the prophet cannot choose more than k random draws. So the optimal value of the following convex program is an upper bound on the payoff of the prophet.

$$\begin{aligned} & \text{maximize} && \sum_i u_i(x_i) \\ & \text{subject to} && \sum_i x_i \leq k && (\tau) \\ & && x_i \geq 0, \quad \forall i \in [n] && (\mu_i) \end{aligned}$$

Define the Lagrangian for the above convex program as

$$L(x, \tau, \mu) = - \sum_i u_i(x_i) + \tau \left(\sum_i x_i - k \right) - \sum_i \mu_i x_i.$$

By KKT stationarity condition, at the optimal assignment, it must be $\frac{\partial}{\partial x_i} L(q, \tau, \mu) = 0$. On the other hand, $\frac{\partial}{\partial x_i} L(q, \tau, \mu) = -F_i^{-1}(1-x_i) + \tau - \mu_i$. Assuming that $x_i > 0$, by complementary slackness $\mu_i = 0$, which then implies that $x_i = 1 - F_i(\tau)$, so $x_i = \Pr[V_i > \tau]$. Furthermore, it is easy to show that the first constraint must be tight, which implies that $\sum_i \Pr[V_i > \tau] = k$. Observe that the contribution of each V_i to the objective value of the convex program is exactly $\mathbf{E}[V_i | V_i > \tau] \Pr[V_i > \tau]$. By using a γ_k -conservative magician we can ensure that each box is opened with probability at least γ_k which implies the contribution of each V_i to the expected payoff of the gambler is $\mathbf{E}[V_i | V_i > \tau] \Pr[V_i > \tau] \gamma_k$ which proves that the expected payoff of the gambler is at least γ_k fraction of optimal objective value of the convex program, which was itself an upper bound on the expected payoff of the prophet. \square

5 Generic Multi Buyer Mechanisms

In this section, we present a formal description of the two generic multi buyer mechanisms outlined toward the end of §3. Throughout the rest of this section we assume that for each buyer $i \in [n]$ we can compute a single buyer mechanism M_i and a corresponding concave benchmark R_i , which together provide α -approximation of the optimal single buyer mechanism/optimal benchmark for buyer i . We show that the resulting multi buyer mechanism will be $\gamma_k \alpha$ -approximation of the optimal multi buyer mechanism in \mathbb{M} , where $k = \min_j k_j$ and γ_k is the optimal magician parameter which is at least $1 - \frac{1}{\sqrt{k+3}}$ (Definition 3).

5.1 Pre-Rounding

This mechanism serves the buyers sequentially (arbitrary order); for each buyer i , it selects a subset S_i of available items and runs the single buyer mechanism $M_i(\bar{x}_i[S_i])$, where \bar{x}_i is an optimal

assignment for the benchmark convex program (\overline{OPT}) , and $\bar{x}_i[S_i]$ denotes the vector resulting from \bar{x}_i by zeroing the entries corresponding to items not in S_i . In particular, this mechanism sometimes precludes some items from early buyers to make them available to late buyers. For each item, the mechanism tries to minimize the probability of preclusion for each buyer by equalizing it for all buyers. Note that, for any given pair of buyer and item, we only care about the probability of preclusion in expectation, where the expectation is taken over the types of other buyers and the random choices of the mechanism. The mechanism is explained in detail in Definition 5.

Definition 5 (γ -Pre-Rounding).

- (I) Solve the convex program (\overline{OPT}) and let \bar{x} be an optimal assignment.
- (II) For each item $j \in [m]$, create an instance of γ -conservative magician (Definition 2) with k_j wands (this will be referred to as the j^{th} magician). We will use these magicians through the rest of the mechanism. Note that γ is a parameter that is given.
- (III) For each buyer $i \in [n]$:
 - (a) For each item $j \in [m]$, write \bar{x}_{ij} on a box and present it to the j^{th} magician. Let S_i be the set of items where the corresponding magicians opened the box.
 - (b) Run $M_i(\bar{x}_i[S_i])$ on buyer i and use its outcome as the final outcome for buyer i .
 - (c) For each item $j \in [m]$, if a unit of item j was allocated to buyer i in the previous step, break the wand of the j^{th} magician.

Note that since \bar{x} is a feasible assignment for convex program (\overline{OPT}) , it must satisfy $\sum_i \bar{x}_{ij} \leq k_j$, so by setting $\gamma \leftarrow \gamma_k$ and by Theorem 4 and Definition 3 we can argue that each S_i includes each item j with probability at least γ_k where γ_k is at least $1 - \frac{1}{\sqrt{k+3}}$.

In order for the above mechanism to retain at least a γ -fraction of the the expected objective value of each $M_i(\bar{x}_i)$, further technical assumptions are needed in addition to $\gamma \leq \gamma_k$. We show that it is enough to assume each R_i has a budget-balanced and cross monotonic cost sharing scheme.

Definition 6 (Budget Balanced Cross Monotonic Cost Sharing Scheme). A function $R : [0, 1]^m \rightarrow \mathbb{R}_+$ has a budget balanced cross monotonic cost sharing scheme iff there exists a cost share function $\xi : [m] \times [0, 1]^m \rightarrow \mathbb{R}_+$ with the following two properties:

- (i) ξ must be budget balanced which means for all $x \in [0, 1]^m$ and $S \subseteq [m]$, $\sum_{j \in S} \xi(j, x[S]) = R(x[S])$.
- (ii) ξ must be cross monotonic which means for all $x \in [0, 1]^m$, $j \in [m]$ and $S, S' \subseteq [m]$, $\xi(j, x[S]) \geq \xi(j, x[S \cup S'])$.

Intuitively, a cost share function associates a fraction of the expected objective value returned by the benchmark function R to each item; and ensures that the fraction associated with each item does not decrease when other items are excluded. In particular, the above assumption holds if $R(x[S])$ is a submodular function of S (e.g., for welfare maximization, assuming that buyers' valuations are submodular¹²). Note that it is enough to show that such a cost sharing function exists; however it is never used in the mechanism and its computation is not required.

¹²We conjecture that it holds in general for revenue maximization, when buyers' valuations are submodular and \mathbb{M} is restricted to mechanisms which use buyer specific item pricing.

Theorem 7 (γ -Pre-Rounding). *Suppose for each buyer i , M_i is an α -approximate incentive compatible single buyer mechanism, and R_i is the corresponding concave benchmark. Also suppose R_i has a budget balanced cross monotonic cost sharing scheme. Then, for any $\gamma \in [0, \gamma_k]$, the γ -pre-rounding mechanism (Definition 5) is dominant strategy incentive compatible (DSIC) mechanism which is in \mathbb{M} and is a $\gamma\alpha$ -approximation of the optimal mechanism in \mathbb{M} .*

Proof. See §A. □

Remark 1. *The γ -pre-rounding mechanism assumes no control and no prior information about the order in which buyers are visited. The order specified in the mechanism is arbitrary and could be replaced by any other ordering which may be unknown in advance. In particular, this mechanism can be adopted to online settings in which buyers are served in an unknown order.*

Corollary 1. *In any setting where Theorem 7 is applicable and when \mathbb{M} includes all feasible BIC mechanisms, the gap between the optimal DSIC mechanism and the optimal BIC mechanism is at most $1/\gamma_k$. This gap is at most 2 (for $k = 1$) and vanishes as $k \rightarrow \infty$. That is because Definition 5 is always DSIC, yet it approximates the optimal mechanism in \mathbb{M} .*

5.2 Post-Rounding

This mechanism runs $M_i(\bar{x}_i)$ simultaneously and independently for all buyers i to compute a tentative allocation/payment for each buyer; it then deallocates some of the items at random to ensure that the supply constraints are met at every instance; it ensures that the probability of deallocation perceived by each buyer (i.e., in expectation over the types of other buyers and random choices of the mechanism) is equalized and therefore simultaneously minimized for all buyers. The payments are also adjusted respectively. The mechanism is explained in detail in Definition 7.

Definition 7 (γ -Post-Rounding).

- (I) *Solve the convex program (\overline{OPT}) and let \bar{x} denote an optimal assignment.*
- (II) *Run $M_i(\bar{x}_i)$ simultaneously and independently for all buyers $i \in [n]$, and let $X'_i \subseteq [m]$ and $P'_i \in \mathbb{R}_+$ denote respectively the allocation (subset of items) and payment computed by $M_i(\bar{x}_i)$ for buyer i .*
- (III) *For each item $j \in [m]$, create an instance of γ -conservative magician (Definition 2) with k_j wands (this will be referred to as the j^{th} magician). We will use these magicians through the rest of the mechanism. Note that γ is a parameter that is given.*
- (IV) *For each buyer $i \in [n]$:*
 - (a) *For each item $j \in [m]$, write \hat{x}_{ij} on a box and present it to the j^{th} magician, where \hat{x}_{ij} is the exact probability¹³ of $M_i(\bar{x}_i)$ allocating a unit of item j to buyer i ; let S_i be the set of items where the corresponding magicians opened the box.*
 - (b) *Let $X_i \leftarrow S_i \cap X'_i$ and $P_i \leftarrow \gamma P'_i$. The final allocation and payment of buyer i is given by X_i and P_i respectively.*
 - (c) *For each item $j \in X_i$, break the wand of the j^{th} magician.*

Note that $\sum_i \hat{x}_{ij} \leq \sum_i \bar{x}_{ij} \leq k_j$; so by setting $\gamma \leftarrow \gamma_k$ and by Theorem 4 and Definition 3 we can argue that each S_i includes each item j with probability at least γ_k where γ_k is at least $1 - \frac{1}{\sqrt{k+3}}$. Consequently, any item that is in X'_i will also be in X_i with probability exactly γ .

¹³Note that \bar{x}_{ij} is only an upper bound on the probability of allocation, so $\hat{x}_{ij} \leq \bar{x}_{ij}$

In order for γ -post-rounding to retain at least a γ -fraction of the the expected objective value of each $M_i\langle\bar{x}_i\rangle$ and preserve incentive compatibility, further technical assumptions are needed in addition to $\gamma \leq \gamma_k$; next, we present a set of assumptions which is sufficient for this purpose¹⁴.

- (A'1) The exact ex ante allocation rule for each $M_i\langle\bar{x}_i\rangle$ (i.e., \hat{x}) must be available (i.e., efficiently computable). Note that \bar{x} is only an upper bound on the ex ante allocation rule.
- (A'2) The objective functions must be of the form $\text{OBJ}_i(t_i, X_i, P_i) = \text{OBJ}_i(t_i, X_i, 0) + c_i P_i$ in which $c_i \in \mathbb{R}_+$ is an arbitrary fixed constant. Also, each $\text{OBJ}_i(t_i, X_i, 0)$ must have cost sharing scheme in X_i which is cross monotonic and budget balanced.
- (A'3) The resulting mechanism must be in \mathbb{M} . In particular, that implies \mathbb{M} may not be restricted to any form of incentive compatibility stronger than Bayesian incentive compatibility (BIC), because the γ -post-rounding is only BIC.
- (A'4) The valuations of each buyer must be in the form of a weighted rank function of some matroid.

Observe that A'2 obviously holds for revenue maximization (because $\text{OBJ}_i(t_i, X_i, P_i) = P_i$), and also for welfare maximization with quasilinear utilities and submodular valuations (because $\text{OBJ}_i(t_i, X_i, P_i) = v_i(t_i, X_i)$ where $v_i(t_i, X_i)$ is the valuation of buyer i for allocation X_i ¹⁵). Next, we formally define A'4.

Definition 8 (Matroid Weighted Rank Valuation). *A valuation function $v : 2^m \rightarrow \mathbb{R}_+$ is a matroid weighted rank valuation iff there exists a matroid $\mathcal{M} = ([m], \mathcal{I})$, and a weight function $w : [m] \rightarrow \mathbb{R}_+$ such that $v(S)$ is equal to the weight of a maximum weight independent subset of S , i.e.,*

$$v(S) = \max_{I \in \mathcal{I} \cap 2^S} \sum_{j \in I} w(j), \quad \forall S \subseteq [m]$$

Matroid weighted rank valuations include additive valuations with demand constraints, unit demand valuations, etc.

Theorem 8 (γ -Post-Rounding). *Suppose for each buyer i , M_i is an α -approximate incentive compatible single buyer mechanism, and R_i is a corresponding concave benchmark. Also suppose the assumptions A'1 through A'4 hold. Then, for any $\gamma \in [0, \gamma_k]$, the γ -post-rounding mechanism (Definition 7) is a Bayesian incentive compatible (BIC) mechanism which is in \mathbb{M} and is a $\gamma\alpha$ -approximation of the optimal mechanism in \mathbb{M} .*

Proof. See §A. □

6 Single Buyer Mechanisms

In this section, we present approximately optimal single buyer mechanisms for several common settings. Each one of the single buyer mechanisms presented in this section satisfies the requirements of one of the generic multi buyer mechanisms of §5, so they can be readily converted to a multi buyer mechanisms. Except for §6.4, we restrict the space of mechanisms to item pricing mechanisms with budget randomization as defined next.

¹⁴I.e., one might come up with other sets of assumptions that are also sufficient.

¹⁵Note that the payment terms cancel out because the utility of the seller is counted toward the social welfare of the mechanism

Definition 9 (Item Pricing with Budget Randomization (IPBR)). *An item pricing mechanism is a possibly randomized mechanism that offers a menu of prices to each buyer and allows each buyer to choose their favorite bundle. The payment of a buyer is equal to the total price of the items in her purchased bundle. Note that the prices offered to different buyers do not need to be identical and buyers can be served sequentially. In the presence of budget constraints, a buyer is allowed to pay a fraction of the price of an item and receive the item with a probability equal to the paid fraction¹⁶. A mechanism is considered an item pricing mechanism if its outcome can be interpreted as such¹⁷.*

Item pricing mechanisms are simple and practical as opposed to optimal BIC mechanisms which often involve lotteries. Also budget randomization allows us to get around the hardness of the knapsack problem faced by the budgeted buyers; in particular, assuming that budgets are large compared to prices, budget randomization can be safely ignored since the optimal integral solution of the knapsack problem approaches its optimal fractional solution.

Table 1 lists several settings for which we obtain a multi buyer mechanism with an improved approximation factor compared to previous best known approximations. For each setting, we present a single buyer mechanism that satisfies the requirements of one of the generic multi buyer mechanisms of §5. The corresponding single buyer mechanisms are presented in detail throughout the rest of this section. Note that the final approximation factor for each multi buyer mechanism is equal to the approximation factor of the corresponding single buyer mechanism multiplied by γ_k ; recall that $\gamma_k \geq 1 - \frac{1}{\sqrt{k+3}}$ which approaches 1 as $k \rightarrow \infty$.

Setting	Approx	Space of Mechanisms	Ref
single item(multi unit), unit demand, budget constraint, revenue maximization	γ_k	item pricing with budget randomization	§6.1
multi item(heterogenous), unit demand, product distribution, revenue maximization	$\frac{1}{2}\gamma_k$	deterministic	§6.2
multi item(heterogenous), additive valuations, product distribution, budget constraint, revenue maximization	$(1 - \frac{1}{e})\gamma_k$	item pricing with budget randomization	§6.3
multi item(heterogenous), additive valuations, correlated distribution with polynomial number of types, budget constraint, matroid constraints, revenue or welfare maximization	γ_k	randomized (BIC)	§6.4

Table 1: Summary of mechanisms obtained using the framework of this paper.

For each single buyer mechanism presented in this paper, the single buyer benchmark function $R(x)$ is defined as the optimal value of some convex program of the following general form, in which u is some concave function, $g_j(\cdot)$ are some convex functions, and \mathbb{Y} is some convex polytope (in the rest of this section we only consider a single buyer, so we will omit the subscript i).

$$\begin{aligned}
& \text{maximize} && u(y) && (\overline{OPT}_1) \\
& \text{subject to} && g_j(y) \leq \bar{x}_j, && \forall j \in [m] \\
& && y \in \mathbb{Y}
\end{aligned}$$

¹⁶A utility maximizing buyer, with submodular valuations and budget constraint, always pays the full price for any item she purchases, except potentially for the last item purchased, for which she must have run out of budget.

¹⁷I.e., an item pricing mechanism may collect all the reports and compute the final outcome along with buyer specific prices, such that the outcome of each buyer would be the same as if each buyer purchased their favorite bundle according to her observed prices, and the prices observed by each buyer should be independent of her report.

Lemma 1. $R(\bar{x})$ is concave, i.e., the optimal value of a convex program of the form (\overline{OPT}_1) is always concave in \bar{x} .

Proof. See section A. □

Note that we can substitute each $R_i(\cdot)$ in the multi buyer benchmark convex program (\overline{OPT}) with the corresponding single buyer benchmark convex program to obtain a combined convex program which can be solved efficiently. If each R_i is captured by a linear program, the combined multi buyer program will also be a linear program.

6.1 Single Item, Unit Demand, Budget Constraint

In this section, we consider a unit-demand buyer with a publicly known budget B and one type of item (i.e., $m = 1$). The only private information of the buyer is her valuation for the item, which is drawn from a publicly known distribution with CDF $F(\cdot)$. To avoid complicating the proofs, we assume that $F(\cdot)$ is continuous and strictly increasing in its domain¹⁸. We present a single buyer mechanism which is optimal among item pricing mechanisms with budget randomization (IPBR). We start by defining the modified CDF function $F^B(\cdot)$ as follows.

$$F^B(v) = \begin{cases} F(v) & v \leq B \\ 1 - (1 - F(v))\frac{B}{v} & v \geq B \end{cases} \quad (F^B)$$

Intuitively, $1 - F^B(p)$ is the probability of allocating the item to the buyer if we offer the item at price p . Note that the buyer only buys if her valuation is more than p which happens with probability $1 - F(p)$; if $p > B$, she will pay her whole budget and only get the item with probability $\frac{B}{p}$, otherwise she pays the full price and receives the item with probability 1. Observe that if we want to allocate the item with probability x we can offer a price of $F^{B(-1)}(1 - x)$ which yields a revenue of $xF^{B(-1)}(1 - x)$ in expectation. Define $\mathfrak{R}(x) = xF^{B(-1)}(1 - x)$ and let $\hat{\mathfrak{R}}(x)$ denote its concave closure (i.e., the smallest concave function that is an upper bound on $\mathfrak{R}(x)$ for every x). We will address the problem of efficiently computing $\hat{\mathfrak{R}}(x)$ later in Lemma 2. Next, we show that the optimal value of the following convex program is equal to the expected revenue of the optimal single buyer IPBR mechanism subject to \bar{x} ; therefore we will define the single buyer benchmark function $R(\bar{x})$ to be equal to the optimal value of this program as a function of \bar{x} .

$$\begin{aligned} & \text{maximize} && \hat{\mathfrak{R}}(x) && (\text{REV}_{\text{single}}) \\ & \text{subject to} && x \leq \bar{x} \\ & && x \geq 0 \end{aligned}$$

Theorem 9. *The revenue of the optimal single buyer IPBR mechanism, subject to an upper bound of \bar{x} on the ex ante allocation rule, is equal to the optimal value of the convex program $(\text{REV}_{\text{single}})$. Furthermore, assuming that x^* is the optimal assignment for the convex program, if $\hat{\mathfrak{R}}(x^*) = \mathfrak{R}(x^*)$, then the optimal mechanism uses a single price $p = F^{B(-1)}(1 - x^*)$ otherwise, it randomized between two prices p^-, p^+ with probabilities θ and $1 - \theta$ for some $\theta \in [0, 1]$ and p^-, p^+ .*

Proof. First, we prove that the expected revenue of the optimal single buyer IPBR mechanism, subject to \bar{x} , is upper bounded by $\hat{\mathfrak{R}}(x^*)$. We then construct a price distribution that obtains

¹⁸The proofs can be modified to work without this assumption.

this revenue. Note that any single buyer IPBR mechanism can be specified as a distribution over prices. Let \mathcal{P} be the optimal price distribution. So the optimal revenue is $\mathbf{E}_{p \sim \mathcal{P}}[p(1 - F^B(p))]$. Note that every price p corresponds to an allocation probability $x = 1 - F^B(p)$. So any probability distribution over p can be specified as a probability distribution over x . Let \mathcal{Q} denote the probability distribution over x that corresponds to price distribution \mathcal{P} , so we can write

$$\begin{aligned} \text{optimal revenue} &= \mathbf{E}_{x \sim \mathcal{Q}} \left[x F^{B(-1)}(1 - x) \right] = \mathbf{E}_{x \sim \mathcal{Q}} [\mathfrak{R}(x)] \\ &\leq \mathbf{E}_{x \sim \mathcal{Q}} [\hat{\mathfrak{R}}(x)] \leq \hat{\mathfrak{R}}(\mathbf{E}_{x \sim \mathcal{Q}}[x]) \end{aligned} \quad \text{By Jensen's inequality}$$

which means the optimal revenue is upper bounded by the value of the convex program for $x = \mathbf{E}_{x \sim \mathcal{Q}}[x]$ ¹⁹; so the optimal revenue is upper bounded by the optimal value of the convex program. That completes the first part of the proof.

Next, we construct an optimal price distribution. If $\hat{\mathfrak{R}}(x^*) = \mathfrak{R}(x^*)$, the optimal price distribution is just a single price $p = F^{B(-1)}(1 - x^*)$; otherwise, by definition of concave closure, there are two points x^- and x^+ and $\theta \in [0, 1]$ such that $x^* = \theta x^- + (1 - \theta)x^+$ and $\hat{\mathfrak{R}}(x^*) = \theta \mathfrak{R}(x^-) + (1 - \theta)\mathfrak{R}(x^+)$. In the latter case, the optimal price distribution offers price $p^- = F^{B(-1)}(1 - x^-)$ with probability θ and offers price $p^+ = F^{B(-1)}(1 - x^+)$ with probability $1 - \theta$. \square

Formally, an optimal single buyer IPBR mechanism can be constructed as follows.

Definition 10 (Mechanism).

- Define the single buyer benchmark $R(\bar{x})$ to be the optimal value of the convex program $(\text{REV}_{\text{single}})$ as a function of \bar{x} .
- Given \bar{x} , solve $(\text{REV}_{\text{single}})$ and let x be an optimal assignment.
- If $\hat{\mathfrak{R}}(x) = \mathfrak{R}(x)$, offer the single price $p = F^{B(-1)}(1 - x)$, otherwise randomize between two prices p^- and p^+ as explained in the proof of Theorem 9.

Theorem 10. *The mechanism of Definition 10 is the optimal revenue maximizing single buyer IPBR mechanism. Furthermore, this mechanism satisfies the requirements of γ -pre-rounding.*

Proof. The proof of the optimality follows from Theorem 9. Furthermore, the benchmark function, $R(x)$, is concave (this follows from Lemma 1) and it has a trivial budget balanced cost sharing scheme (because there is only one item), therefore it meets the requirements of γ -pre-rounding. \square

Next, we address the problem of efficiently computing $\hat{\mathfrak{R}}(\cdot)$.

Lemma 2. *A $(1 + \epsilon)$ -approximation of $\hat{\mathfrak{R}}(\cdot)$, which we denote by $\hat{\mathfrak{R}}_{1+\epsilon}(\cdot)$, can be constructed using a piece-wise linear function with $\ell = \frac{\log L}{\log(1+\epsilon)}$ pieces and in time $O(\ell \log \ell)$ in which L is the ratio of the maximum valuation to minimum non-zero valuation. Note that we need at least $\log_2 L$ bits just to represent such valuations so this construction is polynomial in the input size for any constant ϵ .*

Proof. WLOG, assume that all possible non-zero valuations of the buyer are in the range of $[1, L]$. Let $\ell = \lfloor \frac{\log L}{\log(1+\epsilon)} \rfloor$. For $r = 0 \dots \ell$, consider the prices $p_r = (1 + \epsilon)^{\ell - r}$ and compute the corresponding $x_r = 1 - F^B(p_r)$. Construct $\hat{\mathfrak{R}}_{1+\epsilon}(\cdot)$ by constructing the convex hull of the points:

¹⁹Note that $\mathbf{E}_{x \sim \mathcal{Q}}[x]$ is exactly the probability of allocating the item by the price distribution \mathcal{P} , so it must be no more than \bar{x}

$(0, 0), (x_1, p_1 x_1), (x_2, p_2 x_2), \dots, (x_\ell, p_\ell x_\ell), (1, 0)$. This can be done in time $O(\ell \log \ell)$. Note that $F^{B(-1)}(1 - x)$ is a decreasing function of x so at every $x \in [x_r, x_{r+1}]$, the corresponding price is $F^{B(-1)}(x) \in [p_{r+1}, p_r]$ but $p_r = (1 + \epsilon)p_{r+1}$ therefore at every x , $\mathfrak{R}_{1+\epsilon}(x) \leq \hat{\mathfrak{R}}(x) \leq (1 + \epsilon)\mathfrak{R}_{1+\epsilon}(x)$ which completes the proof. \square

Remark 2. In order to use $\mathfrak{R}_{1+\epsilon}(\cdot)$ in the single buyer mechanism of Definition 10, we need to substitute $(1 + \epsilon)\hat{\mathfrak{R}}_{1+\epsilon}(\cdot)$ in the objective function of the convex program ($\text{REV}_{\text{single}}$) instead of $\hat{\mathfrak{R}}(\cdot)$ for computing the benchmark. Furthermore, the mechanism will be a $(1 - \epsilon)$ -approximation of the optimal single buyer IPBR mechanism. Also notice that finding p^- and p^+ from $\mathfrak{R}_{1+\epsilon}(\cdot)$ is trivial.

6.2 Multi Item (Independent), Unit Demand

In this section, we consider a unit demand buyer with private independent valuations for m items. We assume that for each item j , the buyer's valuation is distributed independently according to a publicly known distribution with CDF $F_j(\cdot)$. We present a single buyer mechanism which is a $\frac{1}{2}$ -approximation of the optimal deterministic revenue maximizing mechanism. To avoid complicating the proofs, we assume that each $F_j(\cdot)$ is continuous and strictly increasing in its domain. Furthermore, we require the distributions to be regular. This mechanism can be used with γ -pre-rounding (Definition 5) to yield a $\frac{1}{2}\gamma_k$ -approximate sequential posted pricing multi buyer mechanism. The previous best approximation mechanism for this setting was a $\frac{1}{6.75}$ -approximate sequential posted pricing mechanism by Chawla et al. (2010)²⁰.

We start by defining $\mathfrak{R}_j(x) = xF_j^{-1}(1 - x)$ for each item j . Because $F_j(\cdot)$ corresponds to a regular distribution, $\mathfrak{R}_j(\cdot)$ is concave as shown in the following lemma.

Lemma 3. If $F(\cdot)$ is the CDF of a regular distribution, the function $\mathfrak{R}(x) = xF^{-1}(1 - x)$ is concave.

Proof. It is enough to show that $\frac{\partial}{\partial x}\mathfrak{R}(x)$ is non-increasing in x . Observe that $\frac{\partial}{\partial x}\mathfrak{R}(x) = F^{-1}(1 - x) - \frac{x}{f(F^{-1}(1 - x))}$ in which $f(\cdot)$ is the derivative of $F(\cdot)$. By substituting $x = 1 - F(p)$, it is enough to show that the resulting function is non-decreasing in p because x is itself non-increasing in p . However, by this substitution we get $\frac{\partial}{\partial x}\mathfrak{R}(x) = p - \frac{1 - F(p)}{f(p)}$ which is non-decreasing in p by definition of regularity. \square

Note that any deterministic mechanism for a unit demand buyer can be interpreted as item pricing. Consequently, $\mathfrak{R}_j(x_j)$ is the maximum revenue that such a mechanism can obtain if item j is allocated with probability x_j . Next, we show that the following convex program gives an upper bound on the expected optimal revenue.

$$\begin{aligned}
& \text{maximize} && \sum_j \mathfrak{R}_j(x_j) && (\text{REV}_{\text{unit}}) \\
& \text{subject to} && x_j \leq \bar{x}_j, && \forall j \in [m] && (\lambda_j) \\
& && \sum_j x_j \leq 1 && (\tau) \\
& && x_j \geq 0, && \forall j \in [m] && (\mu_j)
\end{aligned}$$

²⁰Note that the mechanism of Chawla et al. (2010) does not work for non-regular distributions despite the authors' claim.

Theorem 11. *The revenue of the optimal deterministic single buyer mechanism, subject to an upper bound of \bar{x} on the ex ante allocation rule, is no more than the optimal value of the convex program $(\text{REV}_{\text{unit}})$.*

Proof. Let x^* be the ex ante allocation rule of the optimal single buyer deterministic mechanism. So the expected revenue obtained from each item j is upper bounded by $\mathfrak{R}_j(x_j^*)$ (proof of this claim is essentially the same as the proof of Theorem 9). Consequently, the expected optimal revenue cannot be more than $\sum_j \mathfrak{R}_j(x_j^*)$. Furthermore, the optimal mechanism never allocates more than one item, so $\sum_j x_j^* \leq 1$, and also $x_j^* \leq \bar{x}_j$; therefore x^* is a feasible solution for the convex program; so the expected optimal revenue is upper bounded by the optimal value of the convex program. \square

Next, we present the single buyer mechanism.

Definition 11 (Mechanism).

- Define the benchmark $R(\bar{x})$ to be the optimal value of $(\text{REV}_{\text{unit}})$ as a function of \bar{x} .
- Given \bar{x} , solve $(\text{REV}_{\text{unit}})$ and let x denote an optimal assignment.
- For each item j , assign the price $p_j = F_j^{-1}(1 - x_j)$. WLOG, assume that items are indexed in non-decreasing order of prices, i.e., $p_1 \leq \dots \leq p_m$.
- For each item j , define $r_j = \max(x_j p_j + (1 - x_j) r_{j+1}, r_{j+1})$ and let $r_{m+1} = 0$. Let S^* be the subset of items defined as $S^* = \{j | p_j \geq r_{j+1}\}$.
- Only offer the items in S^* at prices computed in the previous step (i.e., set the price of other items to infinity).

Theorem 12. *The mechanism of Definition 11 obtains at least $\frac{1}{2}$ of the revenue of the optimal deterministic single buyer mechanism in expectation. Furthermore, it satisfies the requirements of γ -pre-rounding.*

Proof. First, we show that this mechanism obtains in expectation at least $\frac{1}{2}$ of its benchmark $R(\bar{x})$, which by Theorem 11 is an upper bound on the optimal revenue. Observe that $R(\bar{x}) = \sum_j x_j p_j$ where x_j is exactly the probability that the valuation of the buyer for item j is at least p_j . Now consider an “adversary replica” who has the exact same valuations as the original buyer, but always buys the item that has the lowest price among all the items priced below her valuation. For any assignment of prices, the revenue obtained from the adversary replica is a lower bound on the revenue obtained from the original buyer. So it is enough to show that the mechanism obtains a revenue of at least $\frac{1}{2} \sum_j x_j p_j$ from the adversary replica. Observe that r_j is exactly the expected revenue obtained from the adversary replica when offered the items in $S^* \cap \{j, \dots, m\}$. In particular, item j is included in S^* if $p_j \geq r_{j+1}$, which implies that the revenue obtained from the purchase of item j , conditioned on purchase, is more than the lower bound on the expected revenue obtained from items $\{j, \dots, m\}$. Finally, observe that the expected revenue obtained from the adversary replica is exactly r_1 . By Lemma 4 we can conclude that $r_1 \geq \frac{1}{2} \sum_j x_j p_j$ which completes the proof of the first claim.

Next, we show that this mechanism satisfies the requirements of γ -pre-rounding. Observe that by Lemma 1, the optimal value of $(\text{REV}_{\text{unit}})$ is a concave function of \bar{x} ; so $R(\bar{x})$ is concave. It only remains to show that $R(\cdot)$ has a budget balanced cross monotonic cost sharing scheme. Let $x_j(\bar{x})$ denote the optimal assignment of variable x_j , in the convex program $(\text{REV}_{\text{unit}})$, as a function of \bar{x} . Define the cost share function

$$\xi(j, \bar{x}) = \mathfrak{R}_j(x_j(\bar{x})).$$

We shall show that ξ is budget balanced and cross monotonic (see Definition 6).

- **Budget balance.** We shall show that for any $\bar{x} \in [0, 1]^m$ and any $S \subseteq [m]$, $R(\bar{x}[S]) = \sum_{j \in S} \xi(j, \bar{x}[S])$. Note that $R(\bar{x}[S]) = \sum_j \mathfrak{R}_j(x_j(\bar{x}[S])) = \sum_{j \in S} \xi(j, x_j(\bar{x}[S]))$ which proved that ξ is budget balanced. Note that $\mathfrak{R}_j(x_j(\bar{x}[S])) = 0$, for any $j \notin S$, because $x_j(\bar{x}[S])$ is forced to be 0.
- **Cross monotonicity.** We shall show that $\xi(j, \bar{x}[S]) \geq \xi(j, \bar{x}[S \cup S'])$, for any $\bar{x} \in [0, 1]^m$ and any $S, S' \subseteq [m]$. Let the Lagrangian of (REV_{unit}) be defined as follows.

$$L(x, \lambda, \tau, \mu) = - \sum_j \mathfrak{R}_j(x_j) + \sum_j \lambda_j (x_j - \bar{x}_j) + \tau (\sum_j x_j - 1) - \sum_j \mu_j x_j$$

The high level idea of the proof is as follows. We show that there is more pressure on the constraint associated with τ when the set of available items is $S \cup S'$ instead of S (i.e., τ is larger for $S \cup S'$); we then show that the optimal x_j can be determined from τ ; in particular, we show that, as the optimal τ increases, the optimal x_j decreases, and consequently $\xi(j, x)$ (which is equal to $\mathfrak{R}_j(x_j)$) decreases as well, which proves ξ is cross monotonic. Next we present the proof in detail.

By KKT stationarity conditions, at the optimal assignment the following holds.

$$\frac{\partial}{\partial x_j} L(x, \lambda, \tau, \mu) = - \frac{\partial}{\partial x_j} \mathfrak{R}_j(x_j) + \lambda_j + \tau - \mu_j = 0$$

First we show that the optimal x_j , and consequently $\xi(j, \bar{x})$, can be determined from the optimal τ ; and they are both non-increasing in τ . Observe that (a) all dual variables must be non-negative, (b) by complementary slackness λ_j may be non-zero only if $x_j = \bar{x}_j$, and (c) complementary slackness implies that μ_j may be non-zero only if $x_j = 0$; therefore, if the optimal τ is given, the optimal assignment for x_j is uniquely²¹ determined by the above equation and the aforementioned complementarity slackness conditions. Let $x_j(\tau)$ denote the optimal assignment of x_j as a function of τ . Due to the concavity of $\mathfrak{R}_j(\cdot)$, and the above KKT condition, we can argue that $x_j(\tau)$ is non-increasing in τ , which also implies that $\xi(j, \bar{x})$ is non-increasing in τ .

Next, we prove by contradiction that ξ is cross monotonic. Let $\tau(\bar{x})$ denote the optimal assignment of τ as a function of \bar{x} . By contradiction, suppose ξ is not cross monotonic, i.e. $\xi(j^*, \bar{x}[S \cup S']) > \xi(j^*, \bar{x}[S])$ for some item j^* ; therefore $\tau(\bar{x}[S]) > \tau(\bar{x}[S \cup S']) \geq 0$. Since $\tau(\bar{x}[S]) > 0$, the inequality associated with τ must be tight (by complementary slackness), so $\sum_j x_j(\tau(\bar{x}[S])) = 1$. On the other hand, for all j , $x_j(\tau(\bar{x}[S \cup S'])) \geq x_j(\tau(\bar{x}[S]))$, with the inequality being strict for $j = j^*$, which means $\sum_j x_j(\tau(\bar{x}[S \cup S'])) > 1$, which is a contradiction. □

Lemma 4. Let p_1, \dots, p_m and x_1, \dots, x_m be two sequences of non-negative real numbers and suppose $\sum_j x_j \leq 1$. For each $j \in [m]$, define $r_j = \max(x_j p_j + (1 - x_j) r_{j+1}, r_{j+1})$ and let $r_{m+1} = 0$. Then $r_1 \geq \frac{1}{2} \sum_j x_j p_j$.

Proof. See section A. □

²¹To avoid complicating the proof, we assume that the functions $\mathfrak{R}_j(\cdot)$ are strictly concave, however this assumption is not necessary.

6.3 Multi Item (Independent), Additive, Budget Constraint

In this section, we consider a buyer with publicly known budget B who has private independent and additive valuations for m items (i.e., her valuation for a bundle of items is the sum of her valuations for individual items in the bundle). We assume the buyer's valuation for each item j is distributed independently according to a publicly known distribution with CDF $F_j(\cdot)$. To avoid complicating the proofs, we assume that each $F_j(\cdot)$ is continuous and strictly increasing in its domain²². We present a single buyer mechanism which is a $(1 - \frac{1}{e})$ -approximation of the optimal revenue maximizing item pricing mechanism with budget randomization (IPBR). This mechanism can be used with γ -pre-rounding (Definition 5) to yield a $(1 - \frac{1}{e})\gamma_k$ -approximate sequential posted pricing multi buyer mechanism. The previous best approximation mechanism for this setting was an $O(1)$ -approximate²³ sequential posted pricing mechanism by Bhattacharya et al. (2010). We should note that the mechanism in Bhattacharya et al. (2010) is more general as it allows the buyers to have demand constraints as well, and it does not allow for budget randomization.

As in §6.1, we start by defining the modified CDF function $F_j^B(\cdot)$ for each item j as follows.

$$F_j^B(v) = \begin{cases} F_j(v) & v \leq B \\ 1 - (1 - F_j(v))\frac{B}{v} & v \geq B \end{cases} \quad (F_j^B)$$

Furthermore, for each item j , let $\mathfrak{R}_j(x) = xF_j^{B-1}(1 - x)$ and let $\widehat{\mathfrak{R}}_j(\cdot)$ be its concave closure as define in §6.1. Also, for each j , define $R_j(\bar{x}_j)$ to be the optimal value of the following convex program as a function of \bar{x}_j .

$$\begin{aligned} & \text{maximize} && \mathfrak{R}_j(x_j) && (\text{REV}_{add}) \\ & \text{subject to} && x_j \leq \bar{x}_j \\ & && x_j \geq 0 \end{aligned}$$

The next theorem provides an upper bound on the revenue of the optimal single buyer IPBR mechanism.

Theorem 13. *The revenue of the optimal single buyer item pricing mechanism with budget randomization (IPBR), subject to an upper bound of \bar{x} on the ex ante allocation rule, is no more than $\min(\sum_j R_j(\bar{x}_j), B)$.*

Proof. For any j , if we were only to sell the item j , by Theorem 9, the maximum revenue we could obtain using an IPBR mechanism would be no more than $R_j(\bar{x}_j)$. Observe that if we compute the optimal price distribution for each item separately, we might only get less revenue because the budget is shared among all items and the buyer might not be able to buy some of the items that she would otherwise buy if there were no other items. That means the actually probability of allocating each item j could be less than the optimal assignment of x_j for the convex program (REV_{add}) ; so the optimal joint price distribution might sell at lower prices; but the extra revenue may only come from lower types which were originally excluded by the optimal single item mechanism. Consequently, the overall revenue from each item j cannot be more than $R_j(\bar{x}_j)$. Finally, observe that the expected revenue of the mechanism cannot be more than B , so it can be no more than $\min(\sum_j R_j(x_j), B)$. \square

²²The proofs can be modified to work without this assumption.

²³ $\frac{1}{96}$

Next, we present $(1 - \frac{1}{e})$ -approximate revenue maximizing single buyer IPBR mechanism.

Definition 12 (Mechanism).

- Define the benchmark $R(\bar{x}) = \min(\sum_j R_j(\bar{x}_j), B)$.
- Given \bar{x} , solve the convex program of $(\text{REV}_{\text{add}})$ for each item j , and let x_j denote an optimal assignment.
- For each item j , if $\hat{\mathfrak{R}}_j(x_j) = \mathfrak{R}_j(x_j)$, offer the single price $p_j = F_j^{B(-1)}(1 - x_j)$, otherwise randomize between two prices p_j^- and p_j^+ with probabilities θ_j and $1 - \theta_j$, as explained in Theorem 9. Note that the randomization must be done for each item independently.

Theorem 14. *The mechanism of Definition 12 obtains at least $1 - \frac{1}{e}$ of the revenue of the optimal single buyer IPBR mechanism. Furthermore, this mechanism satisfies the requirements of γ -pre-rounding.*

Proof. First, we show that the mechanism obtains at least $1 - \frac{1}{e}$ of its benchmark $R(\bar{x})$, which by Theorem 13 is an upper bound on the optimal revenue. Consider an imaginary replica of the buyer who has exactly the same valuations as the original buyer, but has a separate budget B for each item. We call this imaginary buyer the “super replica”. Furthermore, suppose that any payment received from the super replica beyond B is lost (i.e., if the super replica pays Z , the mechanism receives only $\min(Z, B)$). Observe that for any assignment of prices, the payment received from the original buyer and the payment received from the super replica are exactly the same because if the original buyer has’t hit his budget limit then both the original buyer and the super replica will buy the same items and pay the exact same amount. Otherwise, if the original buyer hits his budget limit, the mechanism receives exactly B from both the original buyer and the super replica; therefore we only need to show that the revenue obtained by the mechanism from the super replica is at least $(1 - \frac{1}{e})R(\bar{x})$. Observe that from the view point of the super replica there is no connection between different items, so he makes a decision for each item independently. Let Z_j be the random variable corresponding to the amount paid by the super replica for item j . By Theorem 9, we know that $\mathbf{E}[Z_j] = R_j(\bar{x}_j)$ and the total revenue received by the mechanism is $Z = \min(\sum_j Z_j, B)$. Notice that Z_1, \dots, Z_m are independent random variables in the range of $[0, B]$. By applying Lemma 5, we can argue that $\mathbf{E}[\min(\sum_j Z_j, B)] \geq (1 - \frac{1}{e}) \min(\sum_j \mathbf{E}[Z_j], B) = (1 - \frac{1}{e})R(\bar{x})$ which proves our claim.

Next, we show that the mechanism satisfies the requirements of γ -pre-rounding. Observe that all $R_j(\cdot)$ are concave, and so is $R(\bar{x})$. Furthermore, $R(\bar{x}[S]) = \min(\sum_{j \in S} R_j(\bar{x}_j), B)$ is submodular in S for any $S \subseteq [m]$, and therefore it has a cross monotonic budget balanced cost share scheme (see Definition 6), which completes the proof. \square

Lemma 5. *Let B be an arbitrary positive number and let Z_1, \dots, Z_m be independent random variables such that $Z_j \in [0, B]$, for all j . Then the following inequality holds.*

$$\mathbf{E} \left[\min\left(\sum_j Z_j, B\right) \right] \geq \left(1 - \frac{1}{e^{(\sum_j \mathbf{E}[Z_j])/B}}\right) B \geq \left(1 - \frac{1}{e}\right) \min\left(\sum_j \mathbf{E}[Z_j], B\right)$$

Proof. See section A. \square

6.4 Multi Item (Correlated), Additive, Budget and Matroid Constraints

In this section, we consider a buyer with publicly known budget B who has private correlated additive valuations for m items; furthermore, a bundle of items can be allocated to the buyer only if it is an independent set of a matroid $\mathcal{M} = ([m], \mathcal{I})$, where \mathcal{M} is publicly known; equivalently, instead of treating \mathcal{M} as a constraint on the allocation, we may assume that the buyer has matroid valuations, as defined in Definition 8. We assume that the buyer has a discrete type space T . Let $v_t \in \mathbb{R}_+^m$ denote the buyer's valuation vector corresponding to type $t \in T$, and let $f(t)$ denote its probability. We assume that $f(\cdot)$ is represented explicitly as a part of the input, i.e., by enumerating all types along with their respective probabilities. The only private information of the buyer is her type. We present an optimal single buyer randomized mechanism. This mechanism can be used with γ -post-rounding (Definition 7) to yield a γ_k -approximate multi buyer BIC mechanism. Recall that γ_k is at least $\frac{1}{2}$, and approaches 1 as $k \rightarrow \infty$, which means the resulting multi buyer mechanism approaches the optimal multi buyer mechanism as $k \rightarrow \infty$. Prior to the preliminary version of this paper, the best approximation for this setting was a $\frac{1}{4}$ -approximate BIC mechanism by Bhattacharya et al. (2010)²⁴. At the time of writing the current version, Henzinger and Vidali (2011) has also presented a $\frac{1}{2}$ -approximate BIC mechanism for the same setting. Note that all of the aforementioned mechanisms (including the current paper) have running times polynomial only in $|T|$, which means their running time may not be polynomial in the input size if $|T|$ is of exponential size and $f(\cdot)$ has a compact representation.

Consider the following linear program in which $x_t \in [0, 1]^m$ represents the marginal allocation probabilities for type $t \in T$, and p_t represents the corresponding payment. Also let $r_{\mathcal{M}} : 2^m \rightarrow \{0, \dots, m\}$ denote the rank function of \mathcal{M} . The optimal value of this LP is obviously an upper bound on the optimal revenue.

$$\begin{array}{ll}
\text{maximize} & \sum_{t \in T} f(t) p_t \quad (\text{REV}_{\text{corr}}) \\
\text{subject to} & \sum_{t \in T} f(t) x_{tj} \leq \bar{x}_j, \quad \forall j \in [m] \\
& \sum_{j \in S} x_{tj} \leq r_{\mathcal{M}}(S), \quad \forall t \in T, \forall S \subseteq [m] \\
& v_t \cdot x_t - p_t \geq v_{t'} \cdot x_{t'} - p_{t'}, \quad \forall t, t' \in T \\
& x_t \in [0, 1]^m, \quad \forall t \in T \\
& p_t \in [0, B], \quad \forall t \in T
\end{array}$$

Even though the above LP has exponentially many constraints, it can be solved in polynomial time using the ellipsoid method²⁵. Next, we present a mechanism whose expected revenue is equal

²⁴The mechanism in Bhattacharya et al. (2010) considers demand constraint, which is a special case of matroid constraints.

²⁵See Schrijver (2003) for optimization over matroid polytope.

to the optimal value of the above LP, which also implies that it is optimal.

Definition 13 (Mechanism).

- Define the optimal benchmark $R(\bar{x})$ to be the optimal value of $(\text{REV}_{\text{corr}})$ as a function of \bar{x} .
- Given \bar{x} , solve the LP of $(\text{REV}_{\text{corr}})$ and let x and p be an optimal assignment.
- Let t be the buyer's reported type. Allocate a random subset $X \subseteq [m]$ of items such that X is an independent set of \mathcal{M} and each item $j \in [m]$ is included in X with a marginal probability of exactly x_{tj} . This can be achieved by rounding x_t to a vertex of the matroid polytope using dependent randomized rounding (see Chekuri et al. (2010) and the references therein). Also charge a payment of p_t .

Theorem 15. *The mechanism of Definition 13 is an optimal truthful in expectation revenue maximizing single buyer mechanism, subject to an upper bound of \bar{x} on the ex ante allocation rule. Furthermore, it satisfies all the requirements of the γ -post-rounding.*

Proof. The proof of truthfulness and optimality trivially follows from the linear program of $(\text{REV}_{\text{corr}})$. So, we only focus on proving that this mechanism satisfies the requirements of Theorem 8. First, observe that the benchmark function, $R(\bar{x})$, is concave (this follows from Lemma 1). Second, observe that the matroid constraints can be interpreted as matroid valuations for the buyer. Third, notice that the exact ex ante allocation rule can be readily computed from the LP solution, i.e., $\hat{x}_j = \sum_t f(t)x_{tj}$ is the exact probability of allocating item j . Therefore, the mechanism satisfies the requirements of γ -post-rounding. \square

Remark 3. *Observe that by replacing the objective function of $(\text{REV}_{\text{corr}})$ with $\sum_{t \in T} f(t)v_t \cdot x_t$, we get a truthful in expectation welfare maximizing single buyer mechanism, which can also be used with γ -post-rounding to obtain a γ_k -approximate welfare maximizing BIC multiple buyer mechanism.*

7 Analysis of γ -Conservative Magician

In this section, we present the proof of Theorem 4. We prove the theorem in two parts. In the first part, we show that the thresholds computed by the γ -conservative magician indeed guarantee that each box is opened with an ex-ante probability at least γ , assuming that there is enough wand. In the second part, we show that, for any $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$, the thresholds θ_i are no more than $k - 1$ for all i , which implies that the magician never requires more than k wands. It can be shown that a non-adaptive strategy cannot guarantee a probability of more than $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ for opening each box. Furthermore, in Theorem 5, we show that no algorithm can guarantee a probability of $1 - \frac{1}{\sqrt{2\pi k}} + \epsilon$ or better for opening each box, for any $\epsilon \geq 0$, which implies that $1 - \frac{1}{\sqrt{k+3}}$ is not far from optimal. Below, we repeat the dynamic program for computing s_i^ℓ , θ_i and $F_{w_i}(\cdot)$.

$$s_i^\ell = \Pr[S_i = 1 | W_i = \ell] = \begin{cases} 1 & \ell < \theta_i \\ (\gamma - F_{w_i}(\theta_i - 1)) / (F_{w_i}(\theta_i) - F_{w_i}(\theta_i - 1)) & \ell = \theta_i \\ 0 & \ell > \theta_i \end{cases} \quad (s)$$

$$\theta_i = \min\{\ell | F_{w_i}(\ell) \geq \gamma\} \quad (\theta)$$

$$F_{w_{i+1}}(\ell) = \begin{cases} s_i^\ell x_i F_{w_i}(\ell - 1) + (1 - s_i^\ell x_i) F_{w_i}(\ell) & i \geq 1, \ell \geq 0 \\ 1 & i = 0, \ell \geq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (F_w)$$

Part 1. We show that, assuming there is enough number of wands, the thresholds computed by the dynamic program guarantee that each box is opened with an ex ante probability of at least γ .

- (a) First we prove that $\Pr[W_i \leq \ell] \geq F_{w_i}(\ell)$ by induction on i . The base case is trivial. Suppose the inequality holds for $i \geq 1$, we prove it for $i + 1$ as follows.

$$\begin{aligned}
\Pr[W_{i+1} \leq \ell] &\geq \Pr[W_i \leq \ell - 1] + \Pr[W_i = \ell](1 - s_i^\ell x_i) \\
&= \Pr[W_i \leq \ell - 1] s_i^\ell x_i + \Pr[W_i \leq \ell](1 - s_i^\ell x_i) \\
&\geq F_{w_i}(\ell - 1) s_i^\ell x_i + F_{w_i}(\ell)(1 - s_i^\ell x_i) && \text{by induction hypothesis} \\
&= F_{w_{i+1}}(\ell) && \text{by } (F_w)
\end{aligned} \tag{1.a}$$

Observe that all of the above inequalities are met with equality if every x_i is the exact probability of breaking wand for the corresponding box instead of just an upper bound.

- (b) Next, we show that each box is opened with probability at least γ . We shall show that $\Pr[S_i = 1] \geq \gamma$.

$$\begin{aligned}
\Pr[S_i = 1] &= \sum_{\ell} \Pr[S_i = 1 | W_i = \ell] \Pr[W_i = \ell] \\
&= \sum_{\ell=0}^{\theta_i} s_i^\ell \Pr[W_i = \ell] \\
&= \Pr[W_i < \theta_i] + s_i^{\theta_i} \Pr[W_i = \theta_i] && \text{because } s_i^\ell = 1 \text{ for } \ell < \theta_i \\
&= (1 - s_i^{\theta_i}) \Pr[W_i < \theta_i] + s_i^{\theta_i} \Pr[W_i \leq \theta_i] \\
&\geq (1 - s_i^{\theta_i}) F_{w_i}(\theta_i - 1) + s_i^{\theta_i} F_{w_i}(\theta_i) && \text{by (1.a)} \\
&= F_{w_i}(\theta_i - 1) + s_i^{\theta_i} (F_{w_i}(\theta_i) - F_{w_i}(\theta_i - 1)) \\
&= \gamma && \text{by substituting } s_i^{\theta_i} \text{ from (s)}
\end{aligned}$$

Observe that all of the above inequalities are met with equality if each x_i is the exact probability of breaking a wand for the corresponding box instead of being just an upper bound.

Part 2. We show that when $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$, the γ -conservative magician never requires more than k wands, i.e., we show that $\theta_i \leq k - 1$, for all i . First, we present an interpretation of the magician's dynamic program as a stochastic process on an infinite tape with one unit of infinitely divisible sand.

Definition 14 (Sand Displacement Process). *Consider one unit of infinitely divisible sand which is initially at position 0 on an infinite tape; positions on the tape are labeled from left to right. The sand is gradually moved to the right and distributed over the tape in n rounds. Let $F_{w_i}(\ell)$ denote the total amount of sand in positions $\{0, \dots, \ell\}$ at the beginning of round $i \in [n]$. At each round i the following happens.*

- (I) *The leftmost γ -fraction of the sand is selected by identifying the smallest threshold θ_i such that $F_{w_i}(\theta_i) \geq \gamma$ and then selecting all the sand in positions $\{0, \dots, \theta_i - 1\}$ and selecting a fraction of the sand at position θ_i itself such that the total amount of selected is equal to γ . Formally, if $G_i(\ell)$ denotes the total amount of sand selected from $\{0, \dots, \ell\}$, the selection of*

sand is such that $G_i(\ell) = \min(F_{w_i}(\ell), \gamma)$, for every ℓ . In particular, this implies that only a fraction of the sand at position θ_i itself might be selected, however all the sand to the left of position θ_i is selected.

- (II) The selected sand is moved one position to the right as follows. Simultaneously, for all $\ell \in \{0, \dots, \theta_i\}$, an x_i fraction of the sand selected from position ℓ is moved to position $\ell + 1$. For every ℓ , the total amount of sand in positions $\{0, \dots, \ell\}$ at the end of round i is given by the following equation.

$$F_{w_{i+1}}(\ell) = F_{w_i}(\ell) - x_i G_i(\ell) + x_i G_i(\ell - 1) \quad (F_{eq})$$

It is easy to see that θ_i and $F_{w_i}(\ell)$ resulting from the above process are exactly the same as those computed by the γ -conservative magician.

Consider a conceptual barrier which is at position $\theta_i + 1$ at the beginning of round i and is moved to position $\theta_{i+1} + 1$ for the next round, for each $i \in [n]$. It is easy to verify (i.e., by induction) that the sand never crosses to the right side of the barrier (i.e., $F_{w_{i+1}}(\theta_i + 1) = 1$). The following theorem implies that the sand remains concentrated near the barrier throughout the process.

Theorem 16 (Sand). *Throughout the sand displacement process (Definition 14), at the beginning of each round $i \in [n]$, the following inequality holds.*

$$F_{w_i}(\ell) < \gamma F_{w_i}(\ell + 1) \quad \forall i \in [n], \forall \ell \in [\theta_i] \quad (F_{ineq})$$

Furthermore, the average distance of the sand from the barrier, denoted by d_i , is upper bounded by the following inequalities.

$$d_i \leq \frac{1 - \gamma^{\theta_i + 1}}{1 - \gamma} < \frac{1}{1 - \gamma} \quad \forall i \in [n] \quad (d)$$

The first inequality is strict except for $i = 1$.

Proof. We start by proving inequality (F_{ineq}) , by induction on i . The base case of $i = 1$ is trivial because all the sand is at position 0 and $\theta_1 = 0$. Suppose the inequality holds at the beginning of round i and for all $\ell \in [\theta_i]$, we show that it holds at the beginning of round $i + 1$ and for all $\ell \in [\theta_{i+1}]$. Consider any $\ell \in [\theta_{i+1}]$. Note that $\theta_i \leq \theta_{i+1} \leq \theta_i + 1$, so there are two possible cases:

- **Case 1.** If $\ell \in [\theta_i]$, then

$$\begin{aligned} F_{w_{i+1}}(\ell - 1) &= F_{w_i}(\ell - 1) - x_i G_i(\ell - 1) + x_i G_i(\ell - 2) && \text{by } (F_{eq}). \\ &= (1 - x_i) F_{w_i}(\ell - 1) + x_i F_{w_i}(\ell - 2) && \text{by } G_i(\ell') = F_{w_i}(\ell'), \text{ for } \ell' < \theta_i. \\ &< (1 - x_i) \gamma F_{w_i}(\ell) + x_i \gamma F_{w_i}(\ell - 1) && \text{by induction hypothesis.} \\ &= \gamma (F_{w_i}(\ell) - x_i F_{w_i}(\ell) + x_i G_i(\ell - 1)) && \text{by } G_i(\ell') = F_{w_i}(\ell'), \text{ for } \ell' < \theta_i. \\ &\leq \gamma (F_{w_i}(\ell) - x_i G_i(\ell) + x_i G_i(\ell - 1)) && \text{by } G_i(\ell') \leq F_{w_i}(\ell'), \text{ for all } \ell'. \\ &= \gamma F_{w_{i+1}}(\ell) && \text{by } (F_{eq}). \end{aligned}$$

- **Case 2.** If $\ell = \theta_{i+1} = \theta_i + 1$, then by definition of θ_{i+1} , it must be $F_{w_{i+1}}(\theta_{i+1} - 1) < \gamma$. Furthermore, at the end of round i , all the sand must be in the range $\{0, \dots, \theta_i + 1\}$, so $F_{w_{i+1}}(\theta_{i+1}) = F_{w_{i+1}}(\theta_i + 1) = 1$. Consequently, we can conclude that $F_{w_{i+1}}(\theta_{i+1} - 1) \leq \gamma F_{w_{i+1}}(\theta_{i+1})$ which proves the claim.

Next, we prove that the average distance of the sand from the barrier, d_i , is no more than $\frac{1-\gamma^{\theta_i+1}}{1-\gamma}$ at the beginning of round i .

$$\begin{aligned}
d_i &= \sum_{\ell=0}^{\theta_i} F_{W_i}(\ell) && \text{sand at position } \ell \text{ is counted exactly } \theta_i - \ell + 1 \text{ times} \\
&\leq \sum_{\ell=0}^{\theta_i} \gamma^{\theta_i-\ell} F_{W_i}(\theta_i) && \text{by } (F_{eq}) \\
&\leq \sum_{\ell=0}^{\theta_i} \gamma^{\theta_i-\ell} && \text{because } F_{W_i}(\theta_i) \leq 1 \\
&= \frac{1 - \gamma^{\theta_i+1}}{1 - \gamma}
\end{aligned}$$

Note that the second inequality is strict unless $\theta_i = 0$; furthermore conditioned on $\theta_i = 0$, the third inequality is strict except for round $i = 1$; therefore, the upper-bound is strict except for the first round. \square

Theorem 17 (Barrier). *Consider the process of Definition 14. If $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$ and $\sum_{i=1}^n x_i \leq k$ for some $k \in \mathbb{N}_0$, then the barrier never gets past position k , i.e., $\theta_i \leq k - 1$ for all $i \in [n]$.*

Proof. We shall show that the barrier never gets past position k throughout the process, which proves the theorem (recall that the barrier is defined to be at position $\theta_i + 1$). Let d'_i denote the average distance of the sand from the origin, and let d_i denote the average distance of the sand from the barrier, at the beginning of round i . Observe that $d'_i + d_i = \theta_i + 1$; furthermore, $d'_i = d'_{i-1} + \gamma x_{i-1} = \gamma \sum_{r=1}^{i-1} x_r$, i.e., the average distance of the sand from the origin is increased exactly by γx_{i-1} during round $i - 1$ (because the amount of selected sand is exactly γ and x_{i-1} fraction of the selected sand is moved one position to the right). By applying Theorem 16, we get the following inequality.

$$\theta_i + 1 = d'_i + d_i < \gamma \sum_{r=1}^{i-1} x_r + \frac{1 - \gamma^{\theta_i+1}}{1 - \gamma} \leq \gamma k + \frac{1 - \gamma^{\theta_i+1}}{1 - \gamma}$$

In order to show that the barrier never gets past position k , it is enough to show that the above inequality cannot hold for $\theta_i \geq k$; in fact it is just enough to show that it cannot hold for $\theta_i = k$ ²⁶; alternatively, it is enough to show that the complement of the above inequality holds for $\theta_i = k$.

$$k + 1 \geq \gamma k + \frac{1 - \gamma^{k+1}}{1 - \gamma} \quad (\Gamma)$$

Consider the the stronger inequality $k + 1 \geq \gamma k + \frac{1}{1-\gamma}$ which is quadratic in γ and yields a bound of $\gamma \leq 1 - \frac{1}{1/2 + \sqrt{k+1/4}}$; this bound is in fact imposes a looser constraint than $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$ when $k \geq 7$. Furthermore it can be verified (by direct calculation) that inequality (Γ) holds, for $k < 7$ and $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$. That completes the proof. \square

Theorem 17 implies that a γ -conservative magician requires no more than k wands, assuming that $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$. That completes the proof of Theorem 4.

²⁶Because in order for the barrier to get past position k , it must first fall on position $k + 1$.

8 Multi Unit Demands

In this section, we show that the more general model, in which each buyer may need more than one unit but no more than $\frac{1}{k}$ of all units of each item, can be reduced to the simpler model in which there are at least k units of every item and no buyer demands more than 1 unit of each item.

Definition 15 (Multi Unit Demand Market Transformation). *Let k_j denote the number of units of item j . Define $c_j = \lfloor \frac{k_j}{k} \rfloor$ and divide the units of item j almost equally into c_j bins (i.e., each bin will contain either c_j or $c_j + 1$ units). Create a new item type for each bin (i.e., units from the same bin has the same type, but units from different bins are treated as different types of item).*

Theorem 18. *Let \mathbb{M} be the space of feasible mechanisms, in the original (multi unit demand) market, which do not allocate more than $\frac{1}{k}$ of all units of each item to any single buyer. Similarly, let $\mathbb{M}^{(1)}$ be the space of feasible mechanisms, in the transformed market, which do not allocate more than one unit of each item to any single buyer. Any mechanism in \mathbb{M} can be interpreted as a mechanism in $\mathbb{M}^{(1)}$ and vice-versa with the same allocations/payments. Therefore, in order to find the optimal mechanism in the original market, it is enough to find the optimal mechanism in the transformed market.*

Proof. First, we show that any mechanism in $M \in \mathbb{M}^{(1)}$ can be interpreted as a mechanism in \mathbb{M} . That is trivially true because M allocates to each buyer at most one unit from each bin, which is at most c_j units of each item j of the original market, which is no more than $\frac{1}{k}$ of all units of item j .

Next, we show that any mechanism $M \in \mathbb{M}$ can be interpreted as a mechanism in $\mathbb{M}^{(1)}$. For every j , we create a list L_j of all the bins of item j . L_j is initially sorted in decreasing order of the size of the bins. Let X_{ij}^M be the number of units of item j allocated to buyer i by M . We specify the allocations in the transformed market as follows. For each buyer i we repeat the following, X_{ij}^M times: Allocate one unit from the bin that is first in the list L_j and then move the bin back to the end of the list. It is easy to see that no two units from the same bin are allocated to the same buyer, which completes the proof. \square

Note that by Theorem 18, any mechanism in the original market is equivalent to a mechanism in the transformed market with the exact same allocations/payments from the perspective of buyers. Therefore, WLOG, we can work with the transformed market and only consider mechanisms in this market. However, to use our generic multi buyer mechanisms in the transformed market, the underlying single buyer mechanisms should be capable of handling correlated valuations, because units of the same item, even when labeled with different types, are perfect substitutes from the view point of a buyer. Among the single buyer mechanisms presented in this paper, only the mechanism explained in §6.4 can handle correlated valuations.

9 Conclusion

In this paper, for Bayesian combinatorial auctions, we presented an approximate reduction from multi buyer problem to single buyer problems, which leads to the following conclusions.

- **Market size.** As the ratio of the maximum demand to supply (i.e., $\frac{1}{k}$) decreases, less coordination is required on decisions made for different buyers; i.e., as $\frac{1}{k} \rightarrow 0$, the optimal mechanism treats each buyer almost independently of other buyers. Observe that all of the approximation factors in this paper only depend on k (i.e., γ_k) and not on n . It suggests that,

for characterizing asymptotic properties of such markets, the right parameter to consider is perhaps the ratio of the maximum demand to supply; in particular, notice that the number of buyers is irrelevant.

- **Computational hardness.** For mechanism design problems in a variety of settings, the difficulty of making coordinated optimal decisions for multiple buyers can be avoided by losing a small constant factor in the objective (i.e., losing only a $\frac{1}{\sqrt{k+3}}$ fraction of the objective), therefore the main difficulty of constructing constant factor approximation mechanisms in multi dimensional settings stems from the difficulty of designing single buyer mechanisms, which ultimately stems from the IC constraints in the single buyer problem.

10 Acknowledgment

I thank Jason Hartline for many helpful suggestions and comments. I also thank Azarakhsh Malekian for a fruitful discussion that lead to the proof of the sand theorem (Theorem 16).

References

- Babaioff, M., Lavi, R., and Pavlov, E. (2006). Single-value combinatorial auctions and implementation in undominated strategies. In *SODA*, pages 1054–1063.
- Bhattacharya, S., Goel, G., Gollapudi, S., and Munagala, K. (2010). Budget constrained auctions with heterogeneous items. In *44th ACM Symposium on Theory of Computing*, pages 379–388.
- Blumrosen, L. and Holenstein, T. (2008). Posted prices vs. negotiations: an asymptotic analysis. In *ACM Conference on Electronic Commerce*, page 49.
- Bulow, J. and Roberts, J. (1989). The simple economics of optimal auctions. *Journal of Political Economy*, 97(5):1060–90.
- Chakraborty, T., Even-Dar, E., Guha, S., Mansour, Y., and Muthukrishnan, S. (2010). Approximation schemes for sequential posted pricing in multi-unit auctions. In *WINE*.
- Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. (2010). Multi-parameter mechanism design and sequential posted pricing. In *42nd ACM Symposium on Theory of Computing*, pages 311–320.
- Chawla, S., Malec, D., and Malekian, A. (2011). Bayesian mechanism design for budget-constrained agents. In *ACM Conference on Electronic Commerce*.
- Chekuri, C., Vondrák, J., and Zenklusen, R. (2010). Dependent randomized rounding via exchange properties of combinatorial structures. In *FOCS*, pages 575–584.
- Dhangwatnotai, P., Roughgarden, T., and Yan, Q. (2010). Revenue maximization with a single sample. In *ACM Conference on Electronic Commerce*, pages 129–138.
- Hajiaghayi, M. T., Kleinberg, R. D., and Sandholm, T. (2007). Automated online mechanism design and prophet inequalities. In *AAAI*, pages 58–65.
- Hartline, J. D., Kleinberg, R., and Malekian, A. (2011). Bayesian incentive compatibility and matchings. In *SODA*.

- Hartline, J. D. and Lucier, B. (2010). Bayesian algorithmic mechanism design. In *STOC*, pages 301–310.
- Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*, pages 225–234.
- Henzinger, M. and Vidali, A. (2011). Multi-parameter mechanism design under budget and matroid constraints. In *ESA*, pages 192–202.
- Hill, T. P. and Kertz, R. P. (1992). A survey of prophet inequalities in optimal stopping theory. In *Contemporary Mathematics*.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research*, 6(1):pp. 58–73.
- Schrijver, A. (2003). *Combinatorial Optimization : Polyhedra and Efficiency (Algorithms and Combinatorics)*. Springer.
- Yan, Q. (2011). Mechanism design via correlation gap. In *SODA*.

A Missing Proofs

Proof of Theorem 5. Suppose we create n boxes and in each box, independently, we put \$1 with probability $\frac{k}{n}$. If the magician opens a box containing a \$1, then he gets the \$1 but we break his wand (i.e., $x_i = \frac{k}{n}$). Observe that the expected total prize is k dollars, but because we put a dollar in each box independently, there are some instances in which there are more than k non-empty boxes but the magician cannot win more than k dollars at any instance. Let X_i be the indicator random variable which is 1 iff there is a dollar in box i . The expected total prize is $\mathbf{E}[\sum_i X_i] = k$, but the expected prize that the magician can win is at most $\mathbf{E}[\min(\sum_i X_i, k)]$. It can be verified that $\mathbf{E}[\min(\sum_i X_i, k)] \approx (1 - \frac{k^k}{e^k k!})k$ asymptotically as $n \rightarrow \infty$. In fact, for any positive ϵ , there is a large enough n such that $\mathbf{E}[\min(\sum_i X_i, k)] < (1 - \frac{k^k}{e^k k!} + \epsilon)k$. On the other hand, if a magician can guarantee that every box is opened with probability at least $\gamma = 1 - \frac{k^k}{e^k k!} + \epsilon$, then he will be able to obtain a prize of at least $\sum_i \gamma \mathbf{E}[X_i] = (1 - \frac{k^k}{e^k k!} + \epsilon)k$ in expectation which is a contradiction; therefore it is not possible to make such a guarantee. \square

Proof of Theorem 7. First, we show that each S_i includes each item j with probability at least γ . Observe that for each item j , a sequence of n boxes are presented to the j^{th} magician with probabilities $\bar{x}_{1j}, \dots, \bar{x}_{nj}$ written on them. Since $\sum_i \bar{x}_{ij} \leq k_j$ and $\gamma \in [0, \gamma_k]$, we can argue that each box is opened with probability at least γ (see Theorem 4 and Definition 3); therefore S_i includes each item j with probability at least γ .

Next, we show that the expected objective value of γ -pre-rounding is at least $\gamma\alpha$ -fraction of the expected objective value of the optimal mechanism in \mathbb{M} . Note that by Theorem 2, the expected objective value of the optimal mechanism in \mathbb{M} is upper bounded by the optimal value of (\overline{OPT}) which is $\sum_i R_i(\bar{x}_i)$; therefore, it is enough to show that $\mathbf{E}_{S_i}[R_i(\bar{x}_i[S_i])] \geq \gamma\alpha R_i(\bar{x}_i)$, i.e., the expected objective value that $M_i(\bar{x}_i[S_i])$ obtains from buyer i is at least $\gamma\alpha R_i(\bar{x}_i)$. Let ξ_i be a budget balanced cross monotonic cost share function for $R_i(\cdot)$; then

$$\begin{aligned}
\mathbf{E}_{S_i} [R_i(\bar{x}_i[S_i])] &= \mathbf{E}_{S_i} \left[\sum_{j \in S_i} \xi_i(j, \bar{x}_i[S_i]) \right] && \text{because } \xi_i \text{ is budget balanced} \\
&\geq \mathbf{E}_{S_i} \left[\sum_{j \in S_i} \xi_i(j, \bar{x}_i[\{1, \dots, m\}]) \right] && \text{because } \xi_i \text{ is cross monotonic} \\
&= \sum_{j \in [m]} \Pr[j \in S_i] \xi_i(j, \bar{x}_i) \\
&\geq \sum_{j \in [m]} \gamma \xi_i(j, \bar{x}_i) \\
&= \gamma R_i(\bar{x}_i) && \text{because } \xi_i \text{ is budget balanced}
\end{aligned}$$

Next, we show that the multi buyer mechanism based on γ -pre-rounding is in \mathbb{M} and it is dominant strategy incentive compatible (DSIC). The fact that this mechanism is in \mathbb{M} follows from assumption A6 and the fact that for each item j , the corresponding magician breaks no more than k_j wands, which means no more than k_j units are allocated at any instance. To show that it is DSIC, observe that the only way the reports of other buyers could affect the outcome of buyer i is by affecting S_i , yet $M_i(\bar{x}_i[S_i])$ is a mechanism in \mathbb{M}_i , so it is incentive compatible mechanism for any choice of S_i ; therefore the resulting mechanism is DSIC. Observe that this mechanism also preserves all of the ex post properties of each M_i (e.g., individual rationality). \square

Proof of Theorem 8. First, we show that each S_i includes each item j with probability exactly γ . Observe that for each item j , a sequence of n boxes are presented to the j^{th} magician with probabilities $\hat{x}_{1j}, \dots, \hat{x}_{nj}$ written on them. Since $\gamma \in [0, \gamma_k]$ and $\sum_i \hat{x}_{ij} \leq k_j$ and because each $M_i(\bar{x}_i)$ allocates each item j with probability exactly \hat{x}_{ij} , we can argue that each box is opened with probability exactly γ (see Theorem 4 and Definition 3); therefore S_i includes each item j with probability exactly γ .

Next, we show that γ -post-rounding obtains in expectation at least $\gamma\alpha$ -fraction of the expected objective value of the optimal mechanism in \mathbb{M} . Note that by Theorem 2 the expected objective value of the optimal mechanism in \mathbb{M} is upper bounded by the optimal value of (\overline{OPT}) which is $\sum_i R_i(\bar{x}_i)$; therefore, it is enough to show that $\mathbf{E}_{t_i, X_i, P_i}[\text{OBJ}_i(t_i, X_i, P_i)] \geq \gamma\alpha R_i(\bar{x}_i)$, i.e., the expected objective value that γ -post-rounding obtains from buyer i is at least $\gamma\alpha R_i(\bar{x}_i)$. Let ξ_i be a budget balanced cross monotonic cost share function for OBJ_i as required by A'2; then

$$\begin{aligned}
\mathbf{E}_{t_i, X_i, P_i} [\text{OBJ}_i(t_i, X_i, P_i)] &= \mathbf{E}_{t_i, X_i, P_i} [\text{OBJ}_i(t_i, X_i, 0) + c_i P_i] && \text{By } A'2 \\
&= \mathbf{E}_{t_i, X_i, P_i} \left[\sum_{j \in X_i} \xi_i(j, t_i, X_i) + c_i P_i \right] && \text{because } \xi_i \text{ is budget balanced} \\
&\geq \mathbf{E}_{t_i, X_i, P_i} \left[\sum_{j \in X_i} \xi_i(j, t_i, X'_i) + c_i P_i \right] && \text{because } \xi_i \text{ is cross monotonic} \\
&= \mathbf{E}_{t_i, X'_i, P'_i, S_i} \left[\sum_{j \in X'_i} \Pr[j \in S_i] \xi_i(j, t_i, X'_i) + c_i \gamma P'_i \right] \\
&= \mathbf{E}_{t_i, X'_i, P'_i} \left[\sum_{j \in X'_i} \gamma \xi_i(j, t_i, X'_i) + c_i \gamma P'_i \right] \\
&= \gamma \mathbf{E}_{t_i, X'_i, P'_i} [\text{OBJ}_i(t_i, X'_i, P'_i)] \\
&\geq \gamma \alpha R_i(\bar{x}_i)
\end{aligned}$$

Note that the last step follows because $\mathbf{E}_{t_i, X'_i, P'_i} [\text{OBJ}_i(t_i, X'_i, P'_i)]$ is exactly the expected objective value of $M_i(\bar{x}_i)$ which is at least $\alpha R_i(\bar{x}_i)$.

Next, we show that γ -post-rounding is Bayesian incentive compatible (BIC) and does not over allocate any item. Consider any arbitrary buyer i . Observe that each item $j \in X'_i$ is included in X_i with a probability of exactly γ ; furthermore, by $A'4$, valuations of buyer i can be interpreted as a weighted rank function of some matroid; WLOG, we may assume that X'_i is always an independent set of this matroid²⁷; therefore, the valuation of the buyer for the items in X'_i is additive; consequently, her expected valuation for X_i is exactly γ times her valuation for X'_i . Observe that both the expected valuation and the expected payment of buyer i are scaled by γ for any outcome of $M_i(\bar{x}_i)$ and $M_i(\bar{x}_i)$ itself was incentive compatible; therefore, the resulting mechanism is incentive also incentive compatible. However, the final mechanism is only Bayesian incentive compatible because S_i depends on the types/reports of buyers other than i ²⁸. Also note that the mechanism does not over allocate any item, because for each unit of item j being allocated one of the k_j wands of the j^{th} magician breaks. \square

Proof of Lemma 1. The proof is very similar to the proof of Theorem 1. To show that $R(\bar{x})$ is concave, it is enough to show that for any \bar{x} and \bar{x}' and any $\beta \in [0, 1]$, $R(\beta\bar{x} + (1 - \beta)\bar{x}') \geq \beta R(\bar{x}) + (1 - \beta)R(\bar{x}')$. Let y and y' be the optimal assignments for the convex program subject to \bar{x} and \bar{x}' respectively; then $y'' = \beta y + (1 - \beta)y'$ is also a feasible assignment for the convex program subject to $\beta\bar{x} + (1 - \beta)\bar{x}'$; therefore, $R(\beta\bar{x} + (1 - \beta)\bar{x}')$ must be at least $u(\beta y + (1 - \beta)y')$; on the other hand $u(\cdot)$ is concave, so $u(\beta y + (1 - \beta)y') \geq \beta u(y) + (1 - \beta)u(y') = \beta R(\bar{x}) + (1 - \beta)R(\bar{x}')$. That proves the claim. \square

Proof of Lemma 5. Let $\mu = \sum_j \mathbf{E}[Z_j]$. Define the random variables $Y_j = \max(Y_{j-1} - Z_j, 0)$ and $Y_0 = B$. Observe that for each j , $Y_j = \max(B - \sum_{r=1}^j Z_r, 0)$, so $\min(\sum_{r=1}^j Z_r, B) + Y_j = B$.

²⁷Otherwise, we could replace X'_i by a maximum weight independent subset of X'_i .

²⁸I.e., $\Pr[j \in S_i]$ is equal to γ only in expectation over other buyers' reports

Therefore $\mathbf{E}[\min(\sum_{r=1}^j Z_r, B)] + \mathbf{E}[Y_j] = B$ and to prove the theorem it is enough to show that $\mathbf{E}[Y_m] \leq \frac{1}{e^{\mu/B}} B$. We first show that

$$Y_j \leq (1 - \frac{\mathbf{E}[Z_j]}{B}) Y_{j-1}. \quad (Y_j)$$

Consequently

$$Y_m \leq B \prod_{j=1}^m (1 - \frac{\mathbf{E}[Z_j]}{B}) \quad (1)$$

$$\leq B \frac{1}{e^{\mu/B}} \quad (2)$$

The last inequality follows because $\prod_{j=1}^m (1 - \frac{\mathbf{E}[Z_j]}{B})$ takes its maximum when $\frac{\mathbf{E}[Z_j]}{B} = \frac{\mu}{mB}$ (for all j) and $m \rightarrow \infty$.

To prove the second inequality in the statement of the lemma we can use the fact that $(1 - x^a) \geq (1 - x)a$ for any $a \leq 1$, and conclude that $(1 - \frac{1}{e^{\mu/B}})B \geq (1 - \frac{1}{e^{\min(\mu, B)/B}})B \geq (1 - \frac{1}{e}) \frac{\min(\mu, B)}{B} B = (1 - \frac{1}{e}) \min(\mu, B)$.

To complete the proof, we prove inequality (Y_j) as follows.

$$\begin{aligned} \mathbf{E}[Y_j] &= \mathbf{E}[\max(Y_{j-1} - Z_j, 0)] \\ &\leq \mathbf{E}\left[\max(Y_{j-1} - Z_j \frac{Y_{j-1}}{B}, 0)\right] && \text{because } \frac{Y_{j-1}}{B} \leq 1 \\ &= \mathbf{E}\left[Y_{j-1} - Z_j \frac{Y_{j-1}}{B}\right] && \text{because } \frac{Z_j}{B} \leq 1 \\ &= \mathbf{E}[Y_{j-1}] - \frac{1}{B} \mathbf{E}[Z_j Y_{j-1}] \\ &\leq \mathbf{E}[Y_{j-1}] - \frac{1}{B} \mathbf{E}[Z_j] \mathbf{E}[Y_{j-1}] && \text{because } Z_j \text{ and } Y_{j-1} \text{ are independent.} \\ &= (1 - \frac{\mathbf{E}[Z_j]}{B}) \mathbf{E}[Y_{j-1}] \end{aligned}$$

□

Proof of Lemma 4. To prove the claim, it is enough to show that $\frac{r_1}{\sum_j x_j p_j} \geq \frac{1}{2}$. WLOG, we may assume that $\sum_j p_j x_j = 1$ since we can scale p_1, \dots, p_m by a constant $c = \frac{1}{\sum_j x_j p_j}$ and this will also scale r_1, \dots, r_m by the same constant c , so their ratio is not be affected. Consider the following LP and observe that x_j , p_j , and r_j , as defined in the statement of the lemma, form a feasible assignment for this LP. If we show that the optimal objective value of the LP is bounded below by $\frac{1}{2}$, any feasible assignment yields an objective value of at least $\frac{1}{2}$, and therefore $\frac{r_1}{\sum_j x_j p_j} \geq \frac{1}{2}$ which proves the lemma. In the following LP, p_j and r_j are variables and everything else is constant.

$$\begin{array}{llll}
\text{minimize} & r_1 & & \\
\text{subject to} & r_j \geq x_j p_j + (1 - x_j) r_{j+1}, & \forall j \in [m] & (\alpha_j) \\
& r_j \geq r_{j+1}, & \forall j \in [m] & (\beta_j) \\
& \sum_{j=1}^m x_j p_j \geq 1 & & (\gamma) \\
& p_j \geq 0, & \forall j \in [m] & \\
& r_j \geq 0, & \forall j \in [m+1] &
\end{array}$$

To prove that the optimal value of the above LP is bounded below by $\frac{1}{2}$, we construct a feasible assignment for its dual LP, obtaining a value of $\frac{1}{2}$. The dual LP is as follows.

$$\begin{array}{llll}
\text{maximize} & \gamma & & \\
\text{subject to} & \gamma \leq \alpha_j, & \forall j \in [m] & (p_j) \\
& \alpha_1 + \beta_1 \leq 1 & & (r_1) \\
& \alpha_j + \beta_j \leq (1 - x_{j-1})\alpha_{j-1} + \beta_{j-1}, & \forall j \in \{2, \dots, m\} & (r_j) \\
& 0 \leq (1 - x_m)\alpha_m + \beta_m & & (r_{m+1}) \\
& \alpha_j \geq 0, \beta_j \geq 0, \gamma \geq 0, & \forall j \in [m] &
\end{array}$$

We construct an assignment for the dual LP as follows. Set $\alpha_j = \gamma$ and set $\beta_j = \beta_{j-1} - x_{j-1}\gamma$ for all j , except that for $j = 1$ we set $\beta_1 = 1 - \gamma$. From this assignment we get $\beta_j = 1 - \gamma - \gamma \sum_{\ell=1}^{j-1} x_\ell$. Observe that we get a feasible assignment as long as all β_j resulting from this assignment are non-negative. Furthermore, it is easy to see that $\beta_j \geq 1 - \gamma - \gamma \sum_{\ell=1}^m x_\ell \geq 1 - 2\gamma$ because $\sum_j x_j \leq 1$. Therefore, by setting $\gamma = \frac{1}{2}$, all β_j are non-negative and we always get a feasible assignment for the dual LP with an objective value of $\frac{1}{2}$, which completes the proof. \square